13 Physics selection and HLT performance

13.1 Introduction

A preliminary view of the on-line event-selection scheme and the corresponding physics coverage was presented in the HLT, DAQ and DCS Technical Proposal (TP) [13-3]. Since then the studies have evolved to cope with a new scenario for the start-up of the LHC machine and in response to funding constraints. The LHC start-up scenario has a target to deliver 10 fb^{-1} in one year, now assuming a peak luminosity per fill of 2×10^{33} cm^{-2} s^{-1} which is a factor of two higher than assumed in the TP. At the same time as having to address this higher initial luminosity, financial resources from the HLT/DAQ project had to be re-assigned to meet urgent needs elsewhere in the experiment. As a consequence, the construction of the HLT/DAQ system will have to be staged to match the available funding, so that only a reduced system will be available for the initial running. These changes required a major revision of the Physics and Event Selection Architecture of the experiment, including new ideas for reducing event rates and sizes while retaining as many as possible of the ATLAS physics goals. Needless to say, only the availability of real data will allow this proposal to find a final implementation and a fine tuning of the relative weights of the selection signatures. However, it is important to be able to face this phase with the most complete set of tools and a versatile selection architecture, in order to cope with the obvious unknowns that will likely show up at the time of LHC startup.

As it has been described in Chapter 9, the High Level Trigger (HLT) system of the experiment is composed of two separate event-selection steps, LVL2 and the Event Filter (EF), each with distinctive and complementary features. A common denominator is that they will operate using software algorithms running on commercial computers to test hypotheses of particle identification and apply event-selection criteria. LVL2 will do this with special-purpose algorithms that need to operate in about 10 ms and use only part of the detector information at full granularity, while the EF will have the fully-built events at its disposal and work with an execution time of the order of a second. It is important to maintain a flexible scheme able to adapt easily to changes in machine conditions (e.g. luminosity or background rates). The modularity of the HLT will allow the implementation of different reduction steps at different stages.

A essential input to the HLT process is the seeding of the selection with the results from LVL1. When making performance studies for the HLT, a detailed simulation of the LVL1 result is therefore needed. This identifies the regions of the detector (Regions-of-Interest) where potential candidates for interesting physics objects are found. This simulation, described in Section 13.2, allows for a realistic use of the information coming from LVL1, using the same algorithms and procedures that will be implemented in custom hardware in the experiment.

Given the commonalities between LVL2 and the EF, it was recognized that a coherent and organized approach to the software components of the trigger was needed. The work presented in Section 13.3, which represents an important step forward with respect to the TP, has concentrated on this issue. Common tools have been developed for the event selection, and common data-model components and methods have been identified that can be shared by different algorithms, in particular at LVL2. This approach will ease the implementation of different selection schemes and also simplify the migration of algorithms between trigger levels.

Another important point for new developments has been compliance with the updated detector geometry and with the format of the raw data as it will come from the ReadOut System. This

implies that the algorithms must operate starting from streams of bytes organized according to the readout structure of each detector, in exactly the same way as in the real experiment. This has allowed one to study and understand the implications of converting these byte-streams to the objects needed by algorithms in order to perform trigger selections, including making preliminary measurements of the computing time needed for these conversions.

In Section 13.4 the outcome of the present studies is presented. Particular emphasis has been put on the selection of electrons and photons, both of which stem from an electromagnetic cluster identified at LVL1, and muons. For each of the electron/photon and muon event-selection "vertical slices", a thorough implementation of the approach described above has been attempted. For events that passed the LVL1 selection, bytestream raw data organized according to the actual detector readout format are used by LVL2 algorithms operating within the PESA Steering and Control framework Section 9.5. Simulated RoI information from LVL1 is used to initiate the LVL2 processing. Trigger elements are built using detector information and used to test particle identification hypotheses. For events retained by LVL2, the EF reconstruction and analysis are performed (the EF may or may not be seeded using the LVL2 result) and the final-selection result is made available for off-line use. Results on rejection against the dominant backgrounds and on the efficiencies for typical signals are reported, as well as the rates deriving from each of the selections.

In order to span fully the ATLAS physics coverage, signatures involving jets, hadronic decays of tau leptons, large missing transverse energy and also jets with b-quark content are needed, in addition to the electron, photon and muon ones discussed above; these are also discussed in Section 13.4. As described in Chapter 4, spare on-line resources will be used for B-physics studies, e.g. for luminosities below the peak one.

The global assessment of the event rates to be recorded for off-line analysis, based on the present evaluation for each signature, is made in Section 13.5. A sketch of issues related to the initial phase of the experiment, seen from the selection architecture point of view, is given in Section 13.6.

13.2 The LVL1 trigger simulation

An important ingredient to many HLT tests and studies is the simulation of the LVL1 trigger the result of which serves as input to the HLT trigger processing. The ATLAS LVL1 trigger [13-2] is itself a complex system consisting of the calorimeter trigger, the muon trigger and the Central Trigger Processor (CTP) that makes the final LVL1 event decision. Figure 13-1 gives an overview of the LVL1 trigger; the various components mentioned in the figure will be explained later in this section except for the TTC system (trigger, timing and control) which has no equivalent in the simulation.

The LVL1 trigger simulation is implemented in C++ in the ATLAS offline computing framework Athena. It relies heavily on the ATLAS offline data storage implementation, the so-called Transient Event Store (TES). The structure of the simulation follows closely the structure of the LVL1 trigger hardware. Figure 13-2 shows a package view of the LVL1 simulation. The simulation consists of packages simulating various hardware blocks: the resistive plate chamber (RPC) muon trigger (indicated by the package TrigT1RPC in Figure 13-2), the Muon-to-CTP Interface (MuCTPI, package TrigT1Muctpi), the calorimeter trigger (package TrigT1Calo) and the Central Trigger Processor (package TrigT1CTP). The LVL1 configuration (package TrigT1Config) and the simulation of the Region-of-Interest Builder (package TrigT1RoIB) are provided as addition-

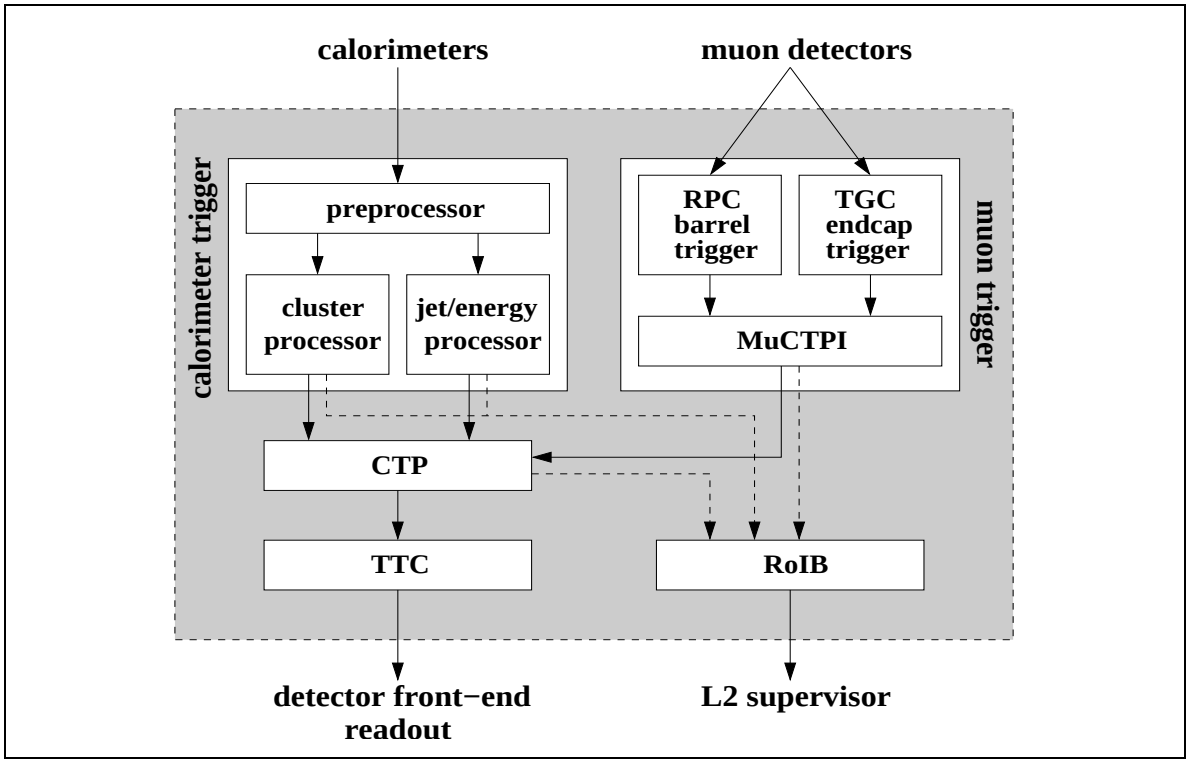

Figure 13-1 An overview of the ATLAS LVL1 trigger system. The Region-of-Interest Builder (RoIB) formally is not a part of the LVL1 trigger. However, it is simulated together with the LVL1 trigger.

al packages. There are also packages for the definition of the LVL1 result raw data object (package TrigT1Result), for classes used by more than one package (package TrigT1Interfaces), and for the conversion of the LVL1 result into the hardware format, the so-called *bytestream* conversion (package TrigT1Result-Bytestream). The various parts of the simulation shown in Figure 13-2 will be explained in the next sections. The simulation of the muon trigger in the endcaps, the signals for which are provided by the thin-gap chambers (package TrigT1TGC), so far exists only as a stand-alone program and will not be discussed in detail in this document.

The interfaces and data formats to be used in the simulation [13-4] were designed to follow as closely as was practical the data formats used in the LVL1 trigger hardware which are documented in [13-5]. Additional information on the LVL1 simulation can be found in [13-6].

13.2.1 Configuration of the LVL1 trigger

The first task of the LVL1 trigger configuration package is to translate the trigger menu, i.e. the collection of event signatures LVL1 is supposed to trigger on, into something that the simulation of the CTP can understand and use in making the event decision based on logical combinations of the inputs delivered by the calorimeter and muon triggers. The LVL1 signatures, or *trigger items*, are combinations of requirements (or *trigger conditions*) on the multiplicities of various kinds of candidate objects found by the calorimeter and muon triggers in the event. (See later subsections for details about the calorimeter and muon triggers and their simulation.)

A simple example of a trigger item is ‘one (or more) electron / photon candidate(s) with transverse energy above 10 GeV and one (or more) muon candidate(s) with transverse momentum above 15 GeV’. In a frequently-used and obvious notation this reduces to ‘1EM10+1MU15’,

where the string ‘EM’ (‘MU’) represents the electron / photon (muon) candidate, and the numbers in front of and behind the string symbolize the required multiplicity and the required transverse momentum, respectively. The combination of a string and a threshold value (like ‘EM10’) is called a *trigger threshold*.

The second task of the LVL1 configuration package is to configure the simulation of the calorimeter and muon triggers to deliver the information required to make the event decision using the trigger menu, i.e. to deliver the multiplicities for the required trigger thresholds. For the example mentioned above, the calorimeter trigger has to be configured to deliver the multiplicity for the threshold ‘EM10’, i.e. the number of electron / photon candidate objects with transverse momentum above 10 GeV, to the CTP. It is obvious that the trigger menu and the trigger thresholds for the calorimeter and muon triggers have to be defined consistently. In particular, all thresholds used in the definition of any trigger condition in any trigger item must be delivered to the CTP by the calorimeter and muon triggers.

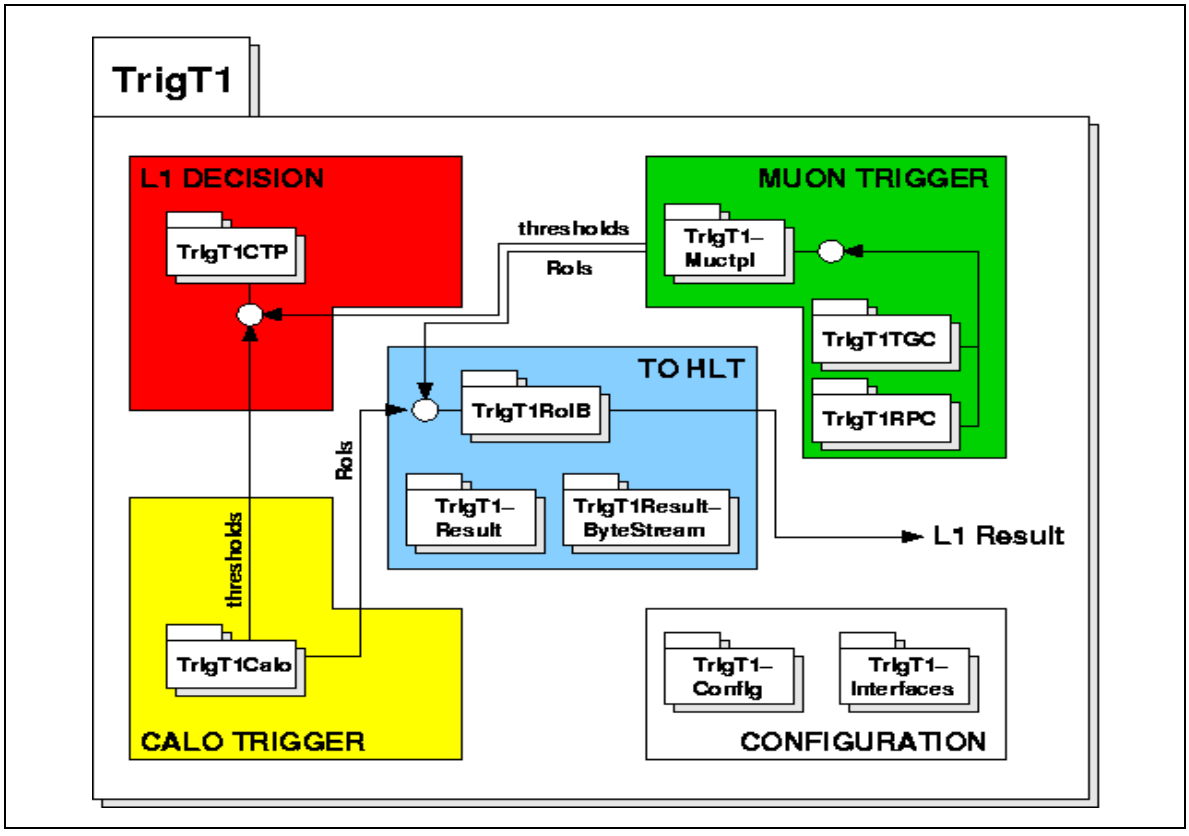

Figure 13-2 A package view of the LVL1 trigger simulation.

Both the trigger menu and the list of required trigger thresholds are defined using XML and are parsed into instances of C++ classes using the Xerces DOM API [13-8]. The parsing of the trigger menu creates an object which contains the information on how the CTP simulation has to discriminate the calorimeter and muon trigger inputs (trigger conditions) and what items have to be built from these conditions. In addition, configuration objects for the calorimeter and muon triggers are created in the configuration process and are stored in the TES for later retrieval by the calorimeter and muon trigger simulations. These objects contain the list of thresholds for which the subsystems have to provide multiplicity information to the CTP simulation.

The LVL1 trigger configuration software is currently being adapted to also be able to configure the LVL1 trigger hardware by deriving the necessary look-up table files and FPGA configuration files from the XML trigger menu and trigger threshold list. Such a common configuration scheme will allow for cross-checks between hardware and software.

13.2.2 The calorimeter trigger and its simulation

The LVL1 calorimeter trigger [13-9] searches for localized energy depositions in the calorimeters due to electrons and photons (electromagnetic clusters), single hadrons and hadronic decays of tau leptons (narrow hadronic clusters) or jets (broader hadronic clusters). For each type of cluster, a number of programmable transverse-energy (E_T) thresholds are provided. With the exception of the jet clusters, isolation requirements can be imposed — these are implemented by applying thresholds on isolation variables associated with the cluster. The multiplicities of the candidate objects of each type are counted for each threshold set, and the multiplicity values are passed on to the CTP to be used in making the LVL1 event decision.

In addition to the local-energy triggers discussed above, the calorimeter trigger calculates global energy sums (total transverse energy and missing transverse energy) which are discriminated against a number of programmable thresholds; the results of the threshold comparisions are also passed to the CTP to be used in making the LVL1 event decision.

The calorimeter trigger uses signals from ~7200 trigger towers which are analogue sums over calorimeter cells in the liquid-argon or scintillator-tile calorimeters. The trigger-tower signals are digitized in the preprocessor electronics, which also performs digital signal processing to calculate transverse energy and make bunch-crossing identification. The resulting data are passed on to two processor subsystems. The Cluster Processor searches for electron/photon and tau/hadron candidates within ~3200 trigger towers of granularity $\Delta\eta\times\Delta\phi = 0.1\times0.1$ in each of the electromagnetic and hadronic calorimeters in the pseudorapidity range $|\eta| < 2.5$. The Jet/Energy Processor searches for jets and makes the global energy sums using coarser (jet) elements of granularity $\Delta\eta\times\Delta\phi = 0.2\times0.2$ over a larger rapidity range ($|\eta| < 3.2$ in case of the jet trigger). Note that it is the preprocessor that sums the trigger towers, independently for the electromagnetic and hadronic calorimeters, to form the larger elements used in the Jet/Energy Processor.

In the electron/photon trigger a candidate object is defined by a local transverse-energy maximum in a region of 2×2 trigger towers in the electromagnetic calorimeter, corresponding to a 0.2×0.2 region in η–ϕ space. Vetos may be applied on the hadronic transverse energy in the corresponding region and/or on the transverse energies in the rings of towers surrounding the 2×2 region in the electromagnetic and hadronic calorimeters. The cluster thresholds are applied on the maximum transverse energy in any edgewise-adjacent pair of electromagnetic-calorimeter trigger towers within the central 2×2 region. (See Refs. [13-2] and [13-10] for a more detailed description of the various calorimeter trigger algorithms.)

In addition to counting multiplicities of candidate objects, Regions-of-Interest (RoIs) are identified and passed, via the Region-of-Interest Builder, to the LVL2 trigger for events that are retained by the LVL1 trigger. The RoIs, which contain information on the geographical location of the object in the calorimeters and on the E_T threshold passed, are used to seed the HLT event-selection process. More detailed information from the LVL1 calorimeter trigger is sent to the DAQ using standard readout links.

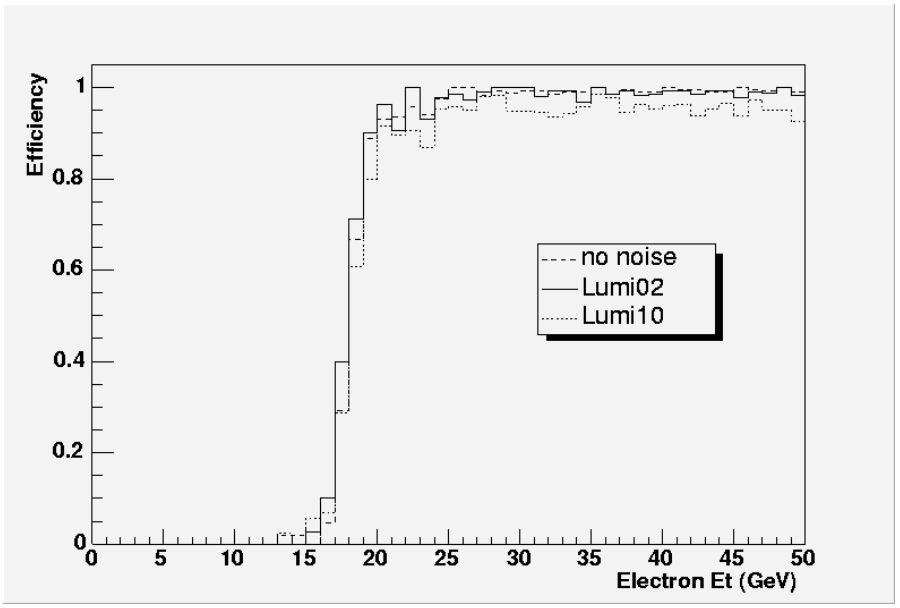

Figure 13-3 Single electron trigger efficiency as a function of the electron transverse energy E_T, using a 17 GeV threshold for different scenarios. See text for details.

The LVL1 calorimeter trigger simulation is designed to reproduce in detail the functionality of the hardware, but it does not entirely duplicate the dataflow. The primary reason for this is efficiency — the hardware trigger will process large amounts of data in parallel, which does not translate well to offline simulation software.

Currently the simulation of the LVL1 calorimeter trigger starts from calorimeter-cell energies which are summed to form tower energies. (A more complete simulation which is in preparation will include details of the signal-processing chain from the calorimeters to the output of the preprocessor, but this is not yet available.) The cell energies, which can be taken from the ATLAS fast simulation or from the detailed GEANT simulation of the calorimeters, are used to build, in a simplified geometrical approach, trigger-tower signals to which calibration and a gaussian noise can be applied. The tower data are passed to the Cluster Processor simulation (electron / photon and tau / hadron finding), and are summed into the coarser jet elements used in the simulation of the Jet / Energy Processor (jet finding and calculation of energy sums).The results from the simulation of the calorimeter trigger are stored in the TES for later retrieval. These include the multiplicity outputs that are used by the CTP in making the event decision, and the details of candidate objects that are used to guide the LVL2 trigger processing (the simulation of the output from the calorimeter trigger simulation to the DAQ has not yet been implemented.).Note that the RoI data from the simulation of the calorimeter trigger are stored in the same bytestream format as will be used in the hardware.

The HLT steering software requires the RoIs to be given with coordinates in η–ϕ space and a value (in GeV) of the E_T threshold that was passed, not in terms of the LVL1 internal data format (which reflects the details of the electronic implementation of the trigger). Software converters are provided to translate the raw 32-bit RoI data words [13-5] into objects, complete with methods to return the required data.

An effort has been made to validate the performance of the calorimeter trigger simulation [13-7]. Here we present some examples of efficiency and rate results for the electron / photon trigger.

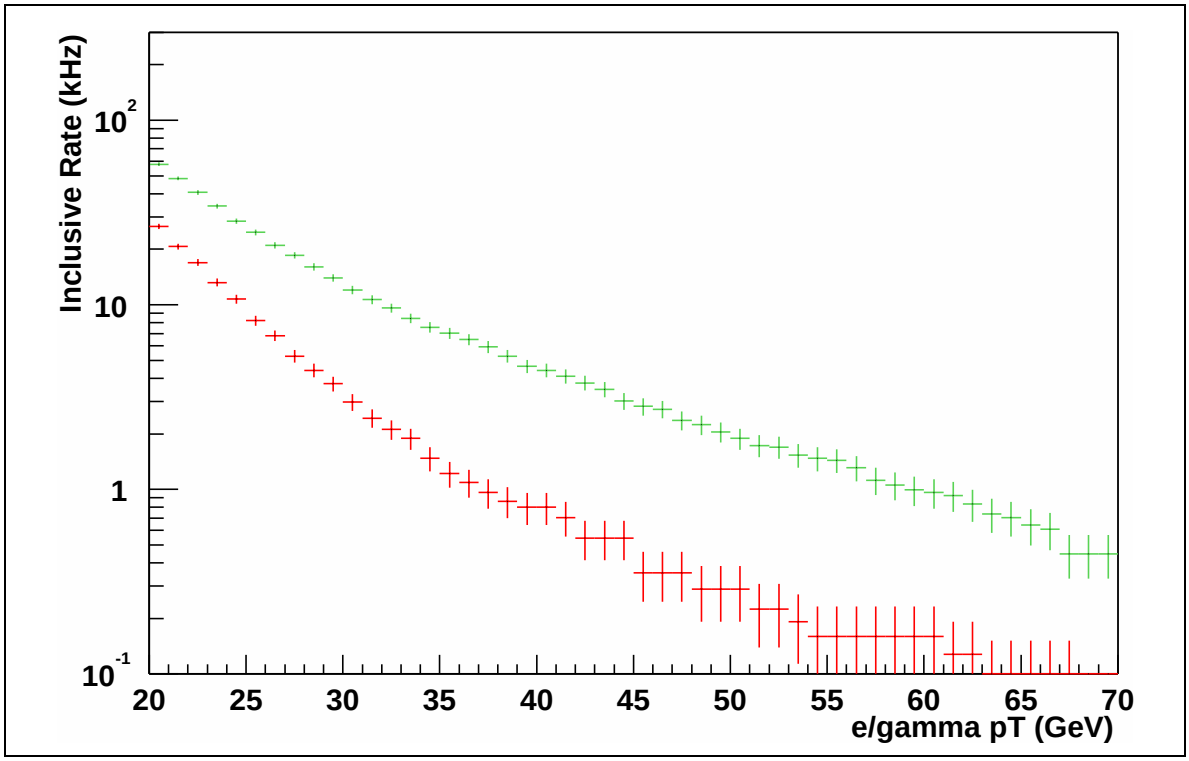

Figure 13-4 Electron/photon trigger rate versus E_T threshold without (top) and with (bottom) isolation requirements at low luminosity 2×10^{33} cm^{-2} s^{-1}. See the text for details.

Figure 13-3 shows the electron trigger efficiency as a function of p_T when requiring that the trigger cluster E_T be greater than 17 GeV. The electrons used in this study came from simulated $Z \rightarrow e^+e^-$ decays, and were required to be well isolated from other high-p_T particles (to ensure that the electron shower was responsible for the trigger) and restricted to a fiducial range $|\eta| < 2.45$. Results are shown without electronic noise or pileup, and also at design luminosity (1×10^{34} cm^{-2} s^{-1}) with electronic noise added. A sharp rise in the efficiency is observed around the threshold value in both cases, with little degradation due to noise and pileup

The rate of the LVL1 electron / photon triggers is dominated by misidentified jets. Figure 13-4 shows the estimated trigger rate for the single electron / photon trigger as a function of transverse-momentum threshold for a luminosity of 2×10^{33} cm^{-2} s^{-1}. The upper and lower bands give the rate before and after applying the isolation requirements. The E_T-threshold scale in the plot is defined so that the efficiency for selecting genuine electrons with transverse-energy equal to the quoted value is 95%. Figure 13-5 presents the rate for the electron / photon-pair trigger for the high-luminosity scenario of 1×10^{34} cm^{-2} s^{-1}, with and without isolation. For this plot, isolation thresholods of 5 GeV, 3 GeV and 2 GeV were used for the electromagentic isolation, hadronic veto and hadronic isolation, respectively.

We have compared the results obtained using the new object-orientated software against those presented in the LVL1 TDR [13-2] which were obtained using the previous FORTRAN-based software. The results are not expected to be identical since the studies used different version of Pythia for the event generation, there have been significant changes to the detector model, and some changes have been made to the RoI-selection and isolation algorithms. Note also that new simulation does not yet include the details of the signal-processing chain for the trigger-tower summation, digitization and bunch-crossing identification. In general there is fair agreement between the old and new results. For example the estimated rates for the single electron / pho-

ton trigger for a luminosity of 2×10^{33} cm^{-2} s^{-1} agree with the earlier simulations to within 20%. In view of the potential sensitivity of the isolation cuts to details of the trigger-tower simulation and to be conservative, we have based our estimates of event and data-rate requirements for the HLT/DAQ system on an extrapolation of the LVL1 TDR results..

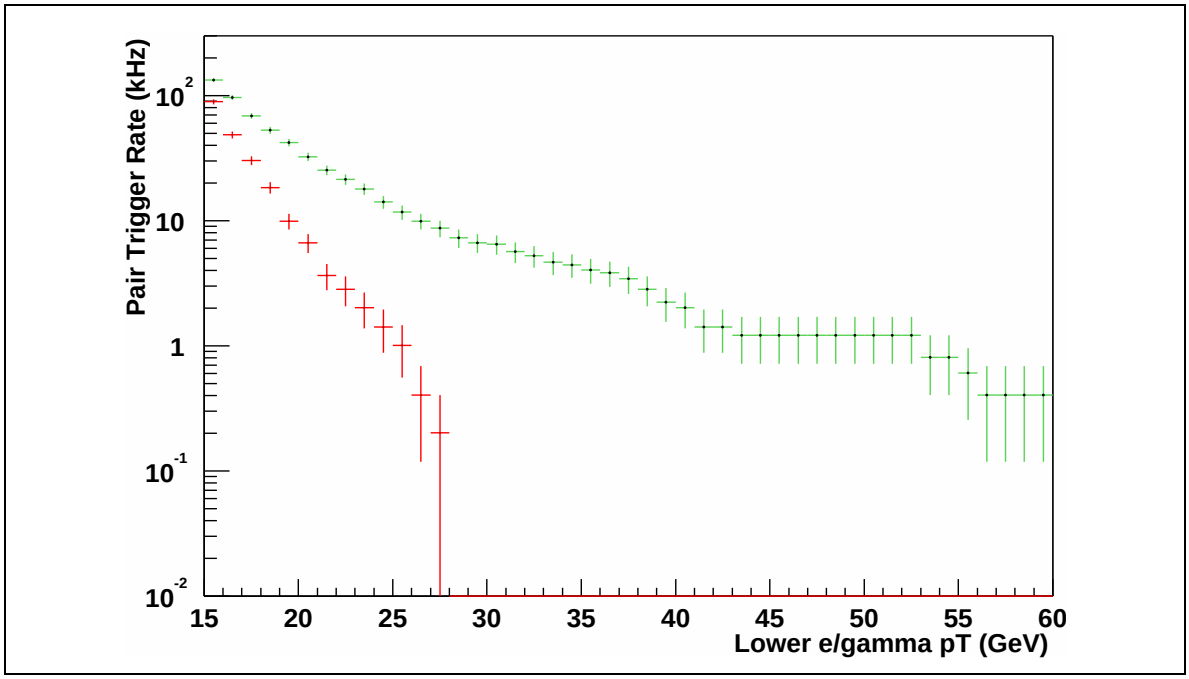

Figure 13-5 Predicted trigger rates for electron pairs without (top) and with (bottom) isolation at high luminosity 1×10^{34} cm^{-2} s^{-1}. See the text for details.

13.2.3 The RPC muon trigger and its simulation

The barrel muon trigger ($|\eta| < 1.05$) uses information from six layers of Resisitive Plate Chamber (RPC) detectors that are organised in three stations [13-11]. The middle RPC station (RPC2) is called the *pivot plane*. The principle of the algorithm that finds muon candidates is as follows [13-12]: each hit found in the pivot plane is extrapolated to the innermost RPC station (RPC1) along a straight line through the innteraction point, and a *coincidence window* is defined around the point of intersection. Since muons are deflected in the magnetic field, the size of the coincidence window defines the p_T threshold of the trigger. A low-p_T muon candidate is found if there is at least one hit in the coincidence window and if in at least one of the stations RPC1 and RPC2 there are hits in both planes of the doublet. This condition must be satisfied for the measurements in the bending plane and also for those in the non-bending plane. If, in addition, there is a hit within a coincidence window in at least one of the two planes of the outermost station RPC3, a high-p_T candidate is found. Again the coincidence condition must be satisfied for the measurements in the bending plane and also for those in the non-bending plane.

In the hardware implementation of the trigger, the alogrithm is implemented using logic that is programmed to recognize valid patterns of hits in the RPCs. All pivot-plane hits are processed in parallel, allowing a trigger with a very short fixed latency. There are three independently programmable thresholds for each of the low-p_T and high-p_T triggers. For each of the 64 sectors of the RPC trigger, up to two muon candidates can be selected and sent to the MuCTPI. If there are

more than two candidates in a sector, the two highest-p_T candidates are retained and a flag indicates that additional candidates were suppressed.

The input to the simulation of the muon trigger logic was provided by a package that simulates the RPC detector system; this was done with the ATLSIM program. The muon detector layout used for this simulation was version “P03” [13-12]. The geometry of the RPC detectors and the stations built from them, as well as the positioning within the muon spectrometer, were reproduced in great detail according to the latest engeering drawings. The material composition and geometry of the single-RPC units were also correctly simulated. The simulation of the RPC detector response was based on results obtained in test-beam experiments. The hits produced in the simulation of charged particles crossing the RPC detectors were collected and stored in output files for use in downstream packages for the simulation of the trigger logic and also for the event reconstruction.

The detector simulation stage is followed by a number of packages which, using Athena algorithms and services, simulate in detail the hardware core and the overall architecture of the LVL1 muon barrel trigger. The hardware simulation is built from a set of classes which reproduce, up to the level of the internal data registers, the behaviour of the basic hardware components: the Coincidence Matrix ASIC (CMA), the Pad Board, the Sector Logic and the Read-Out Driver. The detector data are first accessed by the CMA, in which the RPC digits are translated into bits indicating which input channels fired. Channel masking, time alignment and the introduction of an artificial dead time for the fired channels are possible, although not used yet in the present implementation. The outputs of the eight CMAs belonging to each Pad Board are collected and are searched for valid trigger coincidences in the r–z and r–ϕ views. The Sector Logic then identifies the two highest-p_T muon candidates among all the candidates from all the Pad Boards in the corresponding sector. The output of the Sector Logic, including the addresses and p_T-threshold information for the RoIs, is finally stored in the transient event store of the Athena framework from where it can be retrieved by the MuCTPI simulation package. The CMAs supply information also to the read-out system; this data path is also simulated. The resulting data are organized in a structure that follows exactly the one implemented in the hardware (i.e. bytestream format). This, together with software converters of the interpretation of the bytestream data, allows the use of the RPC data in LVL2 selection algorithms such as muFast ([13-13] and [13-12]).

The architecture model takes care of connecting together the different simulated hardware components and of simulating the data flow in the RPC muon trigger. It is built using a data-driven approach: the event data are arranged in a hierarchical structure going from the RPC digits up to the Sector Logic output. Each level in the hierarchy corresponds to a complete processing step of one hardware component, and the result of each step is available on demand. The “on-demand” approach is very flexible and allows one to save processing time since simulation is only requested for the sectors that contain data.

The architecture simulation makes use of two services that describe the setup of the trigger system: the geometry of the trigger chambers and the cabling of the hardware components. The geometry service accesses the RPC engineering parameters via the AMDB database and builds a model in which the chambers are organized in a suitable way for the LVL1 trigger (i.e. in trigger stations, building a continuous detector surface at fixed radius, and not as appendices of the MDT chambers as it is in the database). The cabling service provides the mapping between the trigger components (RPC readout strips to CMA channels, CMAs to Pad Boards, and Pad Boards to Sector Logic) and holds the configuration of the CMAs. This configuration depends on the muon p_T thresholds which are required. The cabling data as well as the CMA configuration data are read from ASCII files.

The packages that implement the architecture model also provide a fast simulation of the trigger logic that is completely disentangled from the hardware simulation and does not take the timing of signals into account. The main application of this fast simulation is inside the LVL2 selection algorithm muFast, where it is used to emulate the functionality of the LVL1 trigger. Some information on the performance of the RPC muon trigger can be found in [13-13] and [13-14].

13.2.4 The Muon-to-CTP interface and its simulation

The Muon-to-CTP Interface (MuCTPI, see Ref. [13-15]) receives up to two muon candidates from each of the 208 sectors of the barrel (RPC) and endcap (TGC) muon triggers. From these candidates, the multiplicities of muons are calculated for up to six different muon p_T thresholds and sent to the CTP for use in making the trigger decision. If the event is retained by the LVL1 trigger, RoI information on the muon candidates is sent via the Region-of-Interest Builder to the LVL2 trigger (if there are more than 16 candidates, the 16 highest-p_T candidates are retained and a flag is set). More detailed information is sent on a separate link to the DAQ. The data sent to LVL2 and the DAQ conform to the standard ROD bytestream data format. The MuCTPI avoids double-counting of muons which pass through more than one pivot plane, e.g. in the barrel–endcap transition region.

The MuCTPI simulation follows the hardware scheme as closely as possible, down to the data formats used in the hardware. The data flow is emulated using the same stages of processing as in the hardware, including the propagation of error and status bits. Access functions are provided for every type of information available in the hardware. The simulation was originally developed as a stand-alone program for testing the prototype MuCTPI hardware. It has recently been ported to the Athena framework and integrated with the simulation of the RPC trigger on the input side, and with the simulations of the CTP and the RoIB on the output side. The output to the readout is also simulated although this is not yet used within the LVL1 simulation efforts.

13.2.5 The LVL1 CTP and its simulation

The LVL1 trigger event decision is made in the Central Trigger Processor (CTP, see Ref. [13-16]) in the two-step procedure that was discussed above:

The multiplicities of candidate objects from the calorimeter and muon triggers for various p_T thresholds are compared with the *trigger conditions* introduced in Section 13.2.1, checking against simple multiplicity requirements. Depending on the inputs from the calorimeter and muon trigger, each trigger condition takes a logical value TRUE or FALSE.

The trigger conditions (or rather their logical values) are combined using AND, OR and NOT operations to give complex *trigger items*. Each trigger item corresponds to a signature to be triggered on by LVL1 as defined in the trigger menu; gates and prescales can be applied to each individual item. The LVL1 trigger result is the logical OR of all trigger items.

The logical relations between the conditions and the items on one side, and the conditions and the input multiplicities on the other side, are provided by the LVL1 trigger configuration (Section 13.2.1). The CTP provides outputs to the RoIB and to the readout; the information that is sent includes bit patterns for the input signals and for the trigger items before and after prescales and vetos, as well as the final event-decision result.

In the currently-existing prototype hardware implementation of the CTP, the CTP-D ('D' for demonstrator, see Ref. [13-17]), the selection procedure is implemented using two look-up tables (LUT) to compute the trigger conditions and two programmable devices (CPLDs) for the combination of conditions to trigger items. The design of the final CTP is in progress based on the use of large FPGAs.

The current CTP simulation follows closely the CTP-D design — conversion to the final CTP design will be done in due course. First, the input multiplicities from the calorimeter trigger and MuCTPI simulations are collected. The multiplicities that are used in the trigger menu are checked against the respective conditions (the conditions are taken from the C++ object representing the trigger menu that is provided by the configuration step). The conditions are then combined to trigger items in a recursive algorithm. All trigger items are passed through a prescale algorithm, and the logical OR of all items is formed, resulting in the LVL1 event decision (the LVL1 accept or L1A signal which can be expressed as 0 or 1, FALSE or TRUE). The deadtime algorithms that exist in the hardware have not yet been implemented in the simulation. Finally, the CTP result object, which in content and format corresponds precisely to the one provided by the hardware, is formed and stored in the TES for later use by the RoIB.

13.2.6 Interface to the HLT

The interface between the LVL1 trigger and the HLT is the Region-of-Interest Builder (RoIB, see Ref. [13-18]). This device, which is formally part of the DataFlow system, collects the information relevant for LVL2 from the calorimeter and muon triggers and from the CTP. It combines all the data into a single block which it then passes to the LVL2 supervisor assigned to the event in question. The LVL1 data are received in S-LINK format (four links from the calorimeter trigger Cluster Processor, two from the calorimeter Jet/Energy Processor, and one each from the MuCTPI and the CTP). The RoIB has to operate at the highest foreseen LVL1 output rates without introducing additional deadtime.

The RoIB simulation picks up the information stored in the TES by the calorimeter trigger, MuCTPI and CTP simulations, and constructs a LVL1 result raw-data object (RDO). Converters are provided for translating the bytestream format into objects which serve to seed the LVL2 trigger. These contain the value (in GeV) of the threshold that has been passed and the location of the RoI in η–ϕ space.

13.3 Common tools for on-line HLT selection

The HLT algorithms are the basic software components which provide data to derive the trigger decision. These algorithms operate within the context and environment of the PESA Core Software which is discussed from a conceptual design and architectural standpoint in Chapter 9. Section 13.3.1 provides an overview and description from the viewpoint of these HLT algorithms. The objects of a common Event Data Model which algorithms exchange and manipulate are described in Section 13.3.2. An inventory of HLT algorithms intended to operate in the LVL2 environment is given in Section 13.3.3 while those for the EF are described in Section 13.3.4.

13.3.1 Algorithmic View of the Core Software Framework

Unlike their counterparts in the offline Software environment, HLT algorithms must allow themselves to be guided by the PESA Steering, to be seeded by Trigger Elements, and to operate with a restricted set of event data.

To accomplish the Steering guidance of algorithms using Sequence Tables and Trigger Elements, a Seeded approach is required. Trigger Elements characterizing abstract physics objects have a label (e.g., 'electrons' or 'jets') and effectively decouple the Steering and Physics Selection from details of the Event Data Model used by the algorithms. Via the Navigation scheme within the PESA Core Software environment, algorithms may obtain concrete event data associated with a given Trigger Element which define the Seed of restricted and relevant event data fragments upon which they should work.

The Trigger processing itself starts from a LVL1 RoI using predefined sequences of algorithms. These LVL1 RoI objects are associated to Trigger Elements allowing them to be acted upon by the Steering. For each of these Trigger Elements, the Steering executes the required algorithms as defined in a Sequence Table. Hence, it is possible that a given algorithm may be executed *N* times per event. This is fundamentally different than the 'Event Loop' approach of the offline reconstruction paradigm where a given offline algorithm would act only once upon each event.

At LVL2, event data reside within ROBs until actively requested. This allows the LVL2 algorithms to request and process only a small fraction of event data from ROBs, representing a substantial reduction in the network and computation resources required. The first step in this process is the conversion of a geometrical region (*e.g.*, a cone with an extent η and ϕ) into Identifiers; this is accomplished with the HLT RegionSelector.

The HLT RegionSelector [13-20] translates geometrical regions within the fiducial volume of the detector into a set of Identifiers. Presently these Identifiers are `IdentifierHash`es used in the offline software environment. They correspond to elements of appropriate granularity in each sub-detector, usually a DetectorElement. As such, the RegionSelector uses DetectorDescription information during its initialization phase to build an `EtaPhiMap` for each layer (or disk) of a subdetector. This map is essentially a two-dimensional matrix in η and ϕ. Each element consists of a `list` of `IdentifierHash`; the column indices are ϕ floating point numbers while a range (η_{min}, η_{max}) specifies row indices. The input to RegionSelector is the sub-detector under consideration (*i.e.*, Pixel, SCT, TRT, LAr, Tile, MDT, RPC, CSC, or TGC) and the physical extent of the geometrical region. Given the vastly different designs of each subdetector, a subdetector-dependent procedure is used. With knowledge of the layers and/or disks in the region, the RegionSelector searches the $\phi \rightarrow$`IdentifierHash` map which will give a set of `IdentifierHash` is relevant in ϕ region. The last step is to validate each `IdentifierHash` inside the `IdentifierHash`$\rightarrow$ (η_{min}, η_{max}) map.

Interactions with the Data Collection system are hidden from the Algorithm behind a call to StoreGate. Within StoreGate, event data are aggregated into collections within an IdentifiableContainer (IDC) and labelled with an Identifier. Algorithms request event data from StoreGate using the set of Identifiers obtained by the HLT RegionSelector. If the collections are already within StoreGate, it returns them to the HLT algorithm. If not, StoreGate uses the `IOpaqueAddress` to determine which ROBs hold the relevant event data and requests it from the Data Collection system. A ByteStream converter converts the Raw Data into either Raw Data Objects (RDOs) or, by invoking a DataPreparation AlgTool, into Reconstruction Input Objects (RIOs). The obtained RDOs or RIOs are stored within the collections within the IDC within StoreGate.

13.3.2 Event Data Model Components

During 2002 and 2003, there has been a substantial ongoing effort within the HLT, offline, and subdetector communities to establish a common Event Data Model (EDM) between HLT and offline software in the areas of the raw and reconstruction data models. In the discussion that follows, the concept of a DetectorElement is used as an organizing and identifying principle for most EDM objects; these are discussed in Section 13.3.2.1. At the time of writing this document, there has been convergence with respect to the raw data model described in Section 13.3.2.2. Common reconstruction data model classes specific to LVL2 and EF algorithms have been developed and are described in Section 13.3.2.3.

13.3.2.1 Event Data Organization

Event Data (*e.g.*, Raw Data Objects (RDOs) and Reconstruction Input Objects (RIOs)) are aggregated into collections corresponding to adjacent readout channels within the physical detector. These collections reside in an IdentifiableContainer (IDC) with Identifier labels corresponding to the unit of aggregation. For most sub-detectors, the organizing principle is that of the DetectorElement.

In the Pixel detector a DetectorElement is a module, equivalent to a single Silicon wafer; hence there are 1744 Pixel DetectorElements. For the SCT, a DetectorElement is one side of a module, equivalent to a bonded pair of wafers whose strips are oriented in a single direction (*i.e.*, axial or stereo); there are 8176 SCT DetectorElements. For the TRT, a DetectorElement is a planar set of straw tubes representing one row at a constant distance from the module inner wall of straws in a barrel module (*i.e.*, a plane corresponding to the tangential direction in the barrel) and 1/32 in $r\phi$ at a given z of straws in an end-cap wheel; there are 19008 TRT DetectorElements [13-19].

For the calorimeters, the concept of DetectorElement is difficult to define. Instead, the organizing principle for event data is that of the Trigger Tower.

Within the muon spectrometer, for the MDTs, a DetectorElement is a single MDT chamber, where there is at most a single MDT chamber per station, and typically, an MDT chamber has two multilayers. An RPC DetectorElement is the RPC components associated to exactly one barrel muon station; there may be 0, 1 or 2 RPC doublet sets per station and a doublet set may comprise 1, 2 or 4 RPC doublets. A TGC DetectorElement is one TGC η division, or chamber, in a TGC station; there are 24 forward stations in a ring and 48 endcap stations in a ring and there are four rings at each end of the ATLAS detector. Finally, for a CSC DetectorElement is a single CSC chamber, where there is at most a single CSC chamber per station. A CSC chamber typically has two multilayers.

13.3.2.2 Raw Data Model Components

ByteStream Raw Data is ROB-formatted data produced by the ATLAS detector or its simulation [13-19]. It is defined by a set of hierarchical fragments, where only the bottom level, the ROD fragment, is defined by the sub-detector group. The format of the ByteStream has not yet been formally defined. Hence, preliminary “best guesses” have been made as to its structure which may undergo changes in the future.

A Raw Data Object (RDO) is uncalibrated Raw Data converted into an object representing a set of readout channels. Historically this has been referred to as a *Digit*. It is the representation of Raw Data which is put into the Transient Event Store (TES) and is potentially persistifiable.

The purpose of the RDO converters is dual: first a Raw Data ByteStream file can be created by taking the information from the already filled RDOs (in the transient store, from ZEBRA); second, this ByteStream file can then be read back by the converters to fill the RDOs (or the RIOs for LVL2). Since the RDOs are a representation of the specific detector output, its content can change during the life time of the sub-detectors.

A detailed description of the Raw Data Model components is available elsewhere [13-21].

13.3.2.3 Reconstruction Data Model Components

Algorithms interact with Reconstruction Input Objects (RIOs) as opposed to RDOs. For each subdetector system, classes of RIOs have been defined and are described in the following subsections.

13.3.2.3.1 Inner Detector

The implementation of the RIOs makes use of the `IdentifiableContainer` base class, and the collections are also according to the granularity of DetectorElements.

The Pixel and SCT RIOs are Clusters. A Cluster in the Pixel detector is a two-dimensional group of neighbouring readout channels in a DetectorElement. A Cluster in the SCT is a one-dimensional group of neighbouring readout channels in a Detector Element. For Pixel and SCT, there are currently two implementations of the Cluster class: one used for EF and offline and one used for LVL2. The one used at EF has Pixel and SCT sharing the same class. For LVL2 there is a common structure for Pixel, SCT and TRT, but they all have their own concrete classes. For Pixel and SCT there is a base class used for LVL2. There is also an *Extended* class which could potentially be used at EF (which inherits from the LVL2 base class) in the future. Both LVL2 and EF set of cluster classes contain a list of RDO identifiers from which the cluster is built. The number of member functions is limited in both set of classes and the member functions follow the Inner Detector Requirements [13-19]. It is assumed that in the future there will be only one set of RIO classes to be used for LVL2, EF, and offline.

At LVL2, Pixel and SCT RIOs are converted to 3-dimensional coordinates in the ATLAS global coordinate system using the `AlgTools` `SCT_SpacePointTool` and `PixelSpacePointTool`. These tools accept as input a STL `vector` of pointers to Cluster Collections of the appropriate type, `SCT_ClusterCollection` or `PixelClusterCollection`, and return a STL `vector` of objects of the class `TrigSiSpacePoint`. A UML class diagram of the LVL2-specific SpacePoint class `TrigSiSpacePoint` and associated `InDetRecInput` classes is shown in Figure [Ref: fig:spacepoint]

For the Pixels, the creation of SpacePoints consists of combining the local coordinates of Clusters with information on the position and orientation of the DetectorElement to give the global coordinates. The process for the SCT is more complicated since a single SCT detector provides only a one-dimensional measurement. However, an SCT module, consisting of two detectors in a stereo-pair, provide 2-dimensional information. One species of SCT DetectorElement, phi-layer, has strips orientated parallel to the beam axis, the other, *u* or *v* layer, is rotated by ± 40mRad

with respect to the phi-layer DetectorElements. The formation of SpacePoints consists of the following steps:

- Associate each phi-layer Cluster Collection[1] with the corresponding stereo-layer Cluster Collection;
- For each pair of Collections (phi + stereo), take each phi-layer Cluster and search for associated stereo-layer Clusters. If there is more than one associated stereo layer Cluster, a SpacePoint is formed for each (in this case one, at most, will be a correct measurement, the others will form 'ghost' points). If no associated stereo-layer hit is found, a point is created from the phi-layer information alone;
- Calculate the second coordinate (z for the barrel, or R for the end-caps);
- Using information on the position and orientation of the DetectorElement transform to global coordinates.

Note that for the LVL2 SpacePoints some simplifications are made in the interest of speed, as follows:

- No attempt is made to form SpacePoints from Tracks passing close to the edge of a module, where the corresponding stereo-layer Cluster is in a different module.
- Since the stereo and phi layers are separated by a small distance, the trajectory of the track will influence the measurement of the second coordinate. Since the trajectory is not known at the time that SpacePoints are created, there will be a corresponding increase in the uncertainty in the measurement in the second coordinate (R or z).

The TRT RIO is the drift circle of a straw. In the case of the TRT, the same classes are used for LVL2, EF, and offline: those classes are the `DriftCircle` classes part of the set of classes that are also used at LVL2 for Pixel and SCT. The granularity of the TRT RIO is the same as for the RDO: that of a straw, thus the RIO contains an identifier which is the offline identifier for a straw. In the case of the RDO the straw information is uncalibrated and is just the direct content of the detector output, while in the case of the RIO the straw information is calibrated: out of the drift time, a drift radius is obtained. For now, the drift function applied is the same for all straws. In the future the constants that go into the parametrization of this drift function will come from the Interval of Validity Service [13-23].

13.3.2.3.2 Calorimeters

For the Calorimeters, the RIOs are calibrated calorimeter cells (`LArCells` and `TileCells`), imported from the offline reconstruction.

Both `LArCells` and `TileCells` have `CaloCell` as a common base class which represents the basic nature of a observation in the calorimeters an energy, position, time, and quality. A `CaloCell` has been calibrated so that `energy()` returns the physical energy deposit *in the cell* with units of GeV, but without any kind of leakage corrections. Time is given in nanoseconds and refers to when the deposit occurred, relative to the trigger; it should be zero for good hits. Quality reflects how well the input to the system matched the signal model on which the algorithm is based. It is a number with a value between zero to one, giving the significance of the hypothesis that the actual signal is a sampling of the signal model (*i.e.*, it is the integral of a probability dis-

1. There is a Cluster Collection per DetectorElement.

tribution from negative infinity to an observed value of a test statistic and ought to be uniformly distributed between zero and one if the hypothesis is correct).

13.3.2.3.3 Muon Spectrometer

For the barrel Muon Spectrometer it was found expedient to use in part RDOs instead of RIOs as input to the HLT muon selection algorithms. LVL2 uses RDOs, the Event Filter uses RIOs. The RDOs are organized inside the transient event store in identifiable collections and can be accessed in the same way RIOs can be accessed. For the MDTs, RDOs and RIOs are ordered in collections that correspond each to a MDT DetectorElement, i.e. a MDT chamber. Each RDO [13-37],[13-38],[13-39] contains the information of one MDT read-out channels, *i.e.* one MDT tube. The information is the time of the leading edge of the MDT and the charge integrated over the first 20 ns. From this, the uncalibrated drift time can be calculated. A RIO contains the calibrated drift time.

The definition of RPC RDOs is complicated by the fact that the RPCs are trigger chambers. Their read-out is optimised for the trigger task and does not reflect any easily identifiable geometrical structure, as for example an RPC chamber with its strips. Consequently, RPC RDOs are ordered following the read-out structure ordering by PADs and CMAs. Each RDO corresponds to a fired CMA channel [13-40],[13-41]. The RDOs are organized in CMA objects, i.e. in collections corresponding to one CMA each. The CMA objects in turn are organized per PAD, i.e. in collections corresponding to one PAD each. Each PAD contains the information of all coincidences between the inner- and outermost RPC station with the pivot plane within a three-dimensional volume. Neighboring volumes overlap in eta in the inner- and outermost layers, but not in the pivot plane. This allows to assign a unique identifier to each PAD.

One RPC RIO corresponds to an RPC DetectorElement. An RPC RIO contains the information of a collection of RPC strips that fired. There is no simple correspondence between RPC strips and RPC read-out channels. In order to translate a fired RPC read-out channel, which is a CMA read-out channel, into a RPC strip, a cable map and processing of the information of the CMAs for the opposite view is required.

13.3.3 HLT Algorithms for LVL2

13.3.3.1 IDSCAN

IDSCAN (see [13-32] and [13-33]) is a track reconstruction package for LVL2. It takes as input SpacePoints found in the Pixel and SCT Detectors. A series of sub-algorithms (`ZFinder`, `HitFilter`, `GroupCleaner`, `TrackFitter`) then processes these and outputs Tracks and the SpacePoints associated with them.

The `ZFinder` finds the *z*-position of the primary interaction vertex. The algorithm puts all hits into narrow ϕ-bins and extrapolates pairs of hits in each bin back to the beam-line, storing the *z* of intersection in a histogram. It takes as the *z*-position the histogram region with the most entries.

The `HitFilter` finds groups of hits compatible with Tracks from the *z* position found by `ZFinder`. It puts all hits into a histogram binned in ϕ and η. It then finds clusters of hits within

this histogram. It creates a *group* of hits if such a cluster has hits in more than a given number of layers.

The group of hits found by `HitFilter` is used by `GroupCleaner` which splits groups into Tracks and removes noise hits from groups. Each triplet of hits forms a potential track for which p_T, ϕ_0, and d_0 are calculated. It forms groups from these triplets with similar parameters, applying certain quality cuts. It accepts a track candidate if a group contains enough hits.

Finally, the `TrackFitter` verifies track candidates and finds the track parameters by using a standard Kalman-filter-type fitting algorithm adapted from SCTKalman [13-24]. It returns a list of SpacePoints on the Track, the Track parameters, and an error matrix.

13.3.3.2 SiTrack

SiTrack is a track reconstruction package for LVL2 which extends and upgrades a previous algorithm called PixTrig. SiTrack takes Pixel and SCT `SpacePoints` as input and outputs fitted reconstructed Tracks, each storing pointers to the `SpacePoints` used to build it. SiTrack is implemented as a single main algorithm `SiTrack` which instances and executes an user defined list of sub-algorithms (chosen among `STSpacePointSorting`, `STMuonVertex`, `STTrackSeeding`, and `STThreePointFit`).

`STSpacePointSorting` collects pointers to `SpacePoints` coming from the Pixel and SCT detectors and sorts them by module address, storing the result in a Standard Template Library (STL) map. This processing step is performed in order to speed-up data access for the other reconstruction sub-algorithms.

`STMuonVertex` is a primary vertex identification algorithm mostly suitable for low luminosity events with an high p_T muon signature. It is based on track reconstruction inside the LVL1 muon RoI: the most impulsive track is assumed to be the muon candidate and its z impact parameter is taken as the primary vertex position along z.

`STTrackSeeding`, using the sorted `SpacePoint` map and a Monte Carlo Look-Up Table (MC-LUT) linking each B-layer module to the ones belonging to other logical layers, builds track seeds formed by two `SpacePoints` and fits them with a straight line; one or more logical layers can be linked to the B-layer, the latter option being particularly useful if robustness to detector inefficiencies must be improved. If the primary vertex has already been reconstructed by `STMuonVertex`, a fraction of fake track seeds can be rejected during their formation, applying a cut on their z distance from the primary vertex. Otherwise, if no vertex information is available, an histogram whose resolution depends on the number of seeds found is filled with the z impact parameter of each seed; its maximum is then taken as z position for the primary vertex. This vertexing algorithm, which can be operated in both RoI and full scan modes, is best suitable for high luminosity events containing many high p_T tracks (*e.g.*, b-tagging). Independent cuts on r-ϕ and z impact parameters are eventually applied to the reconstructed seeds to further reduce the fake fraction.

`STThreePointFit` extends track seeds with a third `SpacePoint`; it uses a Monte Carlo map associating to each seed a set of module roads[1] the track could have hit passing through the Pixel or SCT detectors. A subset of modules is extracted from each road according to a user defined

1. A road is a list of modules ordered according to the radius at which they are placed starting from the innermost one.

parameter relating to their 'depth' inside it (*e.g.*, the user can decide to use modules at the beginning or in the middle of each road, *etc.*). `SpacePoints` from the selected modules are then used to extend the seed and candidate tracks are fitted with a circle; ambiguities (*e.g.*, tracks sharing at least one SpacePoint) can be solved on the basis of the track quality, leading to an independent set of tracks that can be used for trigger selection or as a seed for further extrapolation.

13.3.3.3 TRTLUT

TRT-LUT is a LVL2 tracking algorithm for track reconstruction in the TRT. It is described in detail elsewhere [13-25]. The algorithm takes as input Hits in the TRT. The algorithmic processing consists of Initial Track Finding, Local Maximum Finding, Track Splitting, and Track Fitting and Final Selection. It outputs the Hits used and Tracks with their parameters.

During the Initial Track Finding, every hit in a three-dimensional image of the TRT detector is allowed to beyond to a number of possible predefined tracks characterized by different parameters. All such tracks are stores in a Look-Up Table (LUT). Every hit increases the probability that a track is a genuine candidate by one unit.

The next step consists of Local Maximum Finding. A two-dimensional histogram is filled with bins in ϕ and $1/p_T$. A histogram for a single track would consists of a "bow-tie" shaped region of bins with entries at a peak in the center of the region. The bin at the peak of the histogram will, in an ideal case, contain all the hits from the Track. The roads corresponding to other filled bins share straws with the peak bin, and thus contain sub-sets of the hits from the track. A histogram for a more complex event would consist of a superposition of entries from individual tracks. Hence, bins containing a complete set of points from each track can be identified as local maxima in the histogram.

The Track Splitting stage of the algorithm analyzes the pattern of hits associated to a track candidate. By rejecting fake candidates composed of hits from several low-p_T tracks, the track splitting step results in an overall reduction by a factor of roughly 2 in the number of track candidates. For roads containing a good track candidate, it identifies and rejects any additional hits from one or more other tracks. The result of the overall Track Splitting step is a candidate that consists of a sub-set of the straws within a road.

The final step of TRT-LUT, Track Fitting and Final Selection, performs a fit in the r-ϕ (z-ϕ) plane for the barrel (end-caps) using a third order polynomial to improve the measurement of ϕ and p_T. Only the straw position is used (*i.e.*, the drift time information is not used). The track is assumed to come from the nominal origin. After the fit, a reconstructed p_T threshold of $0.5\text{GeV}/c$ is applied.

13.3.3.4 TRTKalman

TRT-Kalman [13-26] is a new package based on xKalman++ (see Section [13.3.4.1]). The name is in fact a misnomer since the Kalman filter component of xKalman++ is not used for the TRT; a histogram search and Least Squares fit is used instead.

TRT-Kalman incorporates following modified modules from xKalman:

- `XK_Tracker_TRT`: this reads TRT geometry from ROOT files. It uses `InDetDescr`, `InDetIdentifier` to access necessary Detector Description information;
- `XK_Algorithm`: a strategy is added to perform TRT standalone reconstruction;

- `XK_Track`: a step has been added with fine-tuning of track parameters after the histogramming step and Least Squares fit;
- `XKaTrtMan, XKaTRTRec`: this contains xKalman++ internal steering algorithms;
- `XKaTRTClusters`: this component retrieves `TRT_RDO_Container` from StoreGate filled from a ByteStream file.

13.3.3.5 T2Calo

T2Calo (see Refs. [13-27], [13-28], [13-29], [13-30]) is a clustering algorithm for electromagnetic (EM) showers, seeded by the LVL1 EM trigger RoI positions [13-31]. This algorithm can select isolated EM objects from jets using the cluster E_T and certain shower-shape quantities.

The RIOs are calibrated calorimeter cells (`LArCells` and `TileCells`), imported from the offline reconstruction. Both `LArCells` and `TileCells` have `CaloCell` as common base class. The output (`T2EMCluster`) is a specific LVL2 class containing the cluster energy and position, and the shower-shape variables useful for the selection of EM showers.

The first step in T2Calo is to refine the LVL1 position from the cell with highest energy in the second sampling of the EM calorimeter. This position (η_1, ϕ_1) is later refined in the second sampling by calculating the energy weighted position (η_c, ϕ_c) in a window of 3×7 cells (in $\eta \times \phi$) centred in (η_1, ϕ_1). As described in Ref. [13-28], the steps to perform the jet rejection are the following:

- In sampling 2, $R^{shape}{}_{\eta} = E_{3\times7} \,/\, E_{7\times7}$ is calculated. The expression $E_{n\times m}$ stands for the energy deposited in a window of $n \times m$ around (η_1, ϕ_1).
- In sampling 1, $R^{strip}{}_{\eta} = (E_{1st} - E_{2nd}) \,/\, (E_{1st} + E_{2nd})$ is obtained in a window of $\Delta\eta \times \Delta\phi = 0.125{\times}0.2$ around (η_c, ϕ_c). E_{1st} and E_{2nd} are the energies of the two highest local maxima found, obtained in a strip-by-strip basis. The two ϕ-bins are summed and only the scan in η is considered. A local maximum is defined as a single strip with energy greater than its two adjacent strips.
- The total transverse energy E_T deposited in the EM calorimeter is calculated in a window of 3×7 cells around (η_1, ϕ_1).
- Finally, the energy that leaks into the hadron calorimeter $E^{had}{}_T$ is calculated in a window of size $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ around (η_c, ϕ_c).

13.3.3.6 muFast

The muFast algorithm is a standalone LVL2 tracking algorithm for the Muon Spectrometer. In the past, it existed in the Reference software from ATRIG, and this version is described in detail elsewhere [13-34].

The program is steered by the RoI given by the LVL1 Muon Trigger and uses both RPCs and MDTs measurements. At present this algorithm is limited to the barrel region and it is based on four sequential steps:

1. LVL1 emulation; the muon pattern recognition in the MDT system is initiated by the RPC hits that induced the LVL1 trigger accept. Among these hits, only those related to the pivot plane (middle RPC station) are provided by the muon trigger processor; the ones relat-

ed to the coincidence plane (innermost and outermost RPC stations) have to be identified running a fast algorithm that simulates the basic logic of the LVL1selection.

2. Pattern recognition: it is performed using the RPC hits that induced the LVL1 trigger to define a road in the MDT chambers around the muon trajectory. MDT tubes lying within the road are selected and a contiguity algorithm is applied to remove background hits not associated with the muon trajectory;
3. A straight-line track fit is made to the selected tubes (one per each tube monolayer) within each MDT station. For this procedure the drift-time measurements is used to fully exploit the high measurement accuracy of the muon tracking system. The track sagitta is then evaluated.
4. A fast p_T estimate is made using LUTs. The LUT encodes the linear relationship between the measured sagitta and the Q/p_T, as a function of eta and phi.

The output of this algorithm is the measurement of the muon transverse momentum p_T at the main vertex, eta and phi. The detailed description of the implementation of muFast in the trigger framework is given in [13-63].

13.3.3.7 muComb

The combination of the features of the track measured in the Muon Spectrometer and the Inner Detector (ID) at LVL2 provides a rejection of π and K decays to μ and of fake muons induced by the cavern background. Moreover the combination of the two measurements improves the momentum resolution of reconstructed muons over a large momentum range.

The matching of the Muon Spectrometer tracks and of the ID can be performed extrapolating the ID track to the muon system. The procedure needs to take into account the detector geometry, the material composition and the inhomogeneity of the magnetic field. An accurate extrapolation would require the use of detailed geometry and magnetic field databases, together with a fine tracking. All this would be expensive in terms of CPU time and therefore not acceptable for the LVL2 trigger.

To provide a fast tracking procedure, the effects of the geometry, the materials and of the magnetic field have been described by simple analytic functions of eta and phi. The extrapolation of the ID tracks to the entrance of the Muon Spectrometer is performed using linear extrapolation in two independent projections: the transverse and the longitudinal views. Two coordinates are extrapolated: the z-coordinate and the azimuthal angle phi. The linear extrapolation is corrected using average corrections. In the transverse projection the ID track extrapolation in phi is corrected as follows:

$$\Delta\varphi = \frac{\alpha}{p_T - p_T^0} \qquad \text{13-1}$$

where α is related to the field integral and p_T^0 allows for the transverse energy loss in the material of the calorimeter, that is approximately independent of the track transverse momentum p_T. Both alpha and p_T^0 have been determined by fitting $\Delta\varphi$ of simulated muons as a function of p_T. It is found that $p_T^0 \sim 1.5$, *i.e.* about half of the transverse energy loss of low energy muons, as naively expected. A similar approach has been followed in the case of the extrapolation of the z-coordinate in the longitudinal view.

The matching is done geometrically using cuts on the residuals in each of z and phi. For matching tracks the combined transverse muon momentum is estimated through a weighted average of the independent p_T measurements in the Muon Spectrometer and in the Inner Detector. For each combined track, a χ^2 parameter is used to evaluate the quality of the p_T matching. Thanks to the high quality of the muon p_T measurements in both detectors, secondary muons from π and K decays give typically bad χ^2 matching, and thus can be rejected.

13.3.4 HLT Algorithms for EF

13.3.4.1 xKalman++

xKalman++ is a package for global pattern recognition and Track fitting in the Inner Detector for charged tracks with transverse momentum above 0.5GeV/*c*. A more detailed description of this algorithm is available elsewhere [13-35].

The algorithm starts the track reconstruction in the TRT using a histogramming method or in the Pixel and SCT detector layers using segment search.

The first reconstruction method outputs a set of possible track candidate trajectories defined as an initial helix with a set of parameters and a covariance matrix. As a second step the helix is then used to define a track road through the precision layers, where all the measured clusters are collected. xKalman++ attempts to find all possible helix trajectories within the initial road and with a number of sufficient clusters.

The primary track finding in the Pixels or SCT outputs a set of SpacePoints as an initial trajectory estimation. In the next step these set of space points serve as an input for the Kalman filter-smoother formalism that will add the information from the precision layers. Each reconstructed track is then extrapolated back into the TRT, where a narrow road can be defined around the extrapolation result. All TRT Clusters together with the drift time hits found within this road are then included for the final track-finding and track-fitting steps.

There are three seeding mechanism available in the offline environment: `XKaSeedsAll`, the reconstruction of the full event; `XKaSeedKINE` reconstruction of a region-of-interest and soon available EM calorimeter seeding. In the HLT environment as an EF algorithm xKalman++ will be seeded by the LVL2 result.

After the pattern recognition and Track fitting steps xKalman++ stores the final Track candidates as `SimpleTrack` objects in a `SimpleTrackCollection`. The Track candidate contains the following information:

- Fit procedure used (m-fit or e-fit);
- Helix parameters and their covariance matrix at the end-points of the filter procedure in the precision layers (point on the trajectory closest to the vertex) and in the TRT (point on the trajectory closest to calorimeter);
- Total $\chi 2$ resulting from final fit procedure;
- List of all hits on track from all sub detectors;
- Total number of precision hits *Np*.
- Total number of straw hits *Ns*, empty straws crossed *Ne*, and of drift-time hits *Nt*.

Furthermore, a track candidate is stored in the final output bank if it passes the following cuts:

- The number of precision hits is larger than 5 to 7;
- The ratio $N_S/(N_S+N_e)$ is larger than 0.7 to 0.8;
- The ratio $Nt/N\backslash s$ is larger than 0.5 to 0.7;
- No previously accepted track has the same set of hits as the current one; this last cut removes full *ghost tracks*.

13.3.4.2 iPatRec

iPatRec [13-36] is a pattern recognition algorithm used in the Event Filter that searches for tracks initiated from space-point combinatorial. Pixel and SCT space-points are used to form track-candidates. Candidates are extrapolated to the TRT and drift-time hits added. At the initialization phase, iPatRec creates a geometry data-base describing the properties of each detector module in the precision tracker plus the module's relationship to a simplified material model. This model comprises material "layers" assumed to be either concentric cylinders in the barrel region or planes normal to the beam-axis in the end-cap regions. Additional "layers" represent the TRT detector, beam-pipe and inert support/service material. Track finding, following and fitting procedures make extensive use of this data-base. Another initialization task is to parameters the magnetic field to enable a fast propagation of track parameters between layers.

In the first step of event reconstruction, adjacent raw-data channels are clustered, and space-points produced from these clusters. Each space-point is assigned to one of 7 partitions according to distance from the intersection region. Within each partition the points are ordered according to their azimuthal coordinate. The general procedure is to form track-candidates using space-point combinations from three different partitions subject to criteria on maximum curvature and crude vertex region projectivity. Candidates then undergo a track-fit procedure to give track parameters with covariance at the point of closest approach to the beam-line (perigee parameters). The track follower algorithm propagates these parameters to form an intersect with error ellipse at each "layer" in turn. Clusters are associated to the track from the traversed detectors. Final cluster allocation decisions are taken after a further track-fit. During this fit, energy loss and Coulomb scattering are taken into account by allocating a scattering centre (with associated extra fit parameters) to each "layer" traversed. An active detector region traversed without associated cluster is classified as a 'hole' and retained for material and quality information. Tracks with fit probability greater than 0.001 and a maximum of three holes are extrapolated to the TRT, where a histogramming technique is used to select the TRT hits to be added to the track. Tight cuts are made on the straw residual and on the ratio of found to expected straws, in order to limit high luminosity occupancy effects.

Tracks with cluster(s) in the two innermost Pixel layers plus TRT association are termed primary tracks. Otherwise a maximum of only one hole is allowed: truncated tracks start in the innermost layers but cannot be followed to the outermost layers or TRT; secondary tracks start further out and are required to have TRT association. Various partition combinations are taken to maintain track-finding efficiency for the three types of track even in the event of a higher than expected detector inefficiency. To avoid track duplication, only candidates with two unallocated space-points are initiated, and tracks sharing more than 50% of their clusters are deemed ambiguous whence only the one with higher quality is retained.

To speed up the execution, a preliminary track-finding pass looks only for high quality primary tracks and finishes as soon as one is found. The vertex from this track is then used to subdivide

the space-point partitions into projective slices - greatly reducing the combinatorial load. Impact parameter criteria are adjusted according to the distance to the first cluster to ensure there is no bias against b-, c- or s-particle decays. The code iterates to allow for several initial vertices, very necessary at high luminosity, and reverts to a slower algorithm when no high quality tracks are found. A special fit procedure is available to better handle electron bremsstrahlung. This is invoked from the subsequent combined reconstruction for tracks associated to an EM-calorimeter cluster.

13.3.4.3 LArClusterRec

LArClusterRec is the reconstruction package for electromagnetic clusters in the calorimeter.

In the first step towers are created by summing the cells of the electromagnetic calorimeter and the pre-sampler in depth using a granularity of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ corresponding to the granularity in the second sampling of the EM calorimeter. The input of the tower building are the calibrated calorimeter cells which are produced by the package `LArCellRec`.

In the next step a sliding window algorithm is used. In case a local maximum is found with a total energy in the window above a given transverse energy threshold, clusters are created which are subsequently stored in the cluster container. To reconstruct the cluster energy and position is calculated in a given window.[1] The cluster energy is corrected for η and ϕ modulations and leakage outside the cluster in a given window. In the region between the barrel and end-cap calorimeter the cluster energy is in addition corrected for energy losses using the energy deposit in the crack scintillators. The η position in the first and second sampling is corrected for s-shapes, which is a geometrical effect. The ϕ position is corrected for an offset, which is also a geometry effect.

13.3.4.4 egammaRec

EgammaRec is designed to calculate useful quantities to separate clusters in the electromagnetic calorimeter from jets. To do so, electromagnetic cluster information as well as tracking information is used.

In the electromagnetic calorimeter electrons are narrow objects, while jets tend to have a broader profile. Hence, shower shapes can be used to reject jets. This is handled by the `EMShowerBuilder` which calls different algorithms which calculate diverse quantities using the information in the first and second sampling of the electromagnetic calorimeter as well as the leakage into the first sampling of the hadronic calorimeter.

Cluster and track information is combined in the `TrackMatchBuilder`. For a given cluster all tracks are examined in the given window around the cluster position. In case more than one track is found, the one with the highest p_T is retained. If the *E/p* ratio is $0.5 < E/p < 1.5$, the track match is successful. In the subsequent particle identification step the information provided by egammaRec can be used. In the case of an electron hypothesis, jets can be rejected by analysis of the EM shower shapes, tight track quality cuts, *E/p*, and the position match in η and ϕ between the cluster and the tracks. Photons can be selected by analysing the EM shower shapes, reconstruction of conversions in the Inner Detector, and possibly a track veto for non-converted photons.

1. This window can be different from the one used for the sliding window algorithm.

13.3.4.5 Moore

Moore is an offline track reconstruction package for the Muon Spectrometer. A detailed description of Moore is available elsewhere [13-42]. Moore reconstructs track in the full η range (barrel+endcaps), but we will restrict the description to the barrel region of the Muon spectrometer since at present only this region is considered from the trigger chain (Muon vertical slice).

In the standard offline configuration Moore takes as input collections of digits or clusters inside the Muon Spectrometer and outputs fitted reconstructed tracks whose parameters are expressed at the first measured point inside the Muon Spectrometer. The reconstruction flow is desegregated into sequential steps, and each step is driven by an Algorithm module (`MooMakeXXX`) that builds partial or final reconstruction objects. Each algorithm retrieves object created by the previous modules from StoreGate and builds transient objects that are subsequently recorded in StoreGate where they become available for others algorithms. Data and algorithms are strictly separated: the algorithms should know the structure of the data objects that they are accessing or producing, but the objects are not dependent from the algorithms. The existence of only a dependence from algorithms to data and the flowing sequence of the reconstruction steps allows to establish which algorithm will produce an object at run-time.

Inside the MDTs the drift distance is calculated from the drift time, by applying various corrections: such as the TOF, the second coordinate, the propagation along the wire, the Lorenz effect. From the 4 tangential lines the best one is found. All the MDT segments of the outer station are combined with those of the Middle layer. The MDT hits of each combination are added to the phi-hits of the ϕ Segment, forming outer track candidates. All the successfully fitted candidates are kept for further processing (`MooMakeRoads`).

The successful outer track is subsequently used to associate inner station MDT hits. A final track is defined as a successfully fitted collection of trigger hits and MDT hits from at least two layers (`MooMakeTracks`). The parameters of the fitted tracks are referred to the first measured point inside the MuonSpectrometer. In order to be used for physics studies an extrapolation to the vertex of the track parameters is needed. To accomplish this task a different offline package, MuonIdentification (MuId), should be used. {MuonIdentification} accesses the Moore-track and propagate it to the vertex, parameterizing the multiple scattering as scattering planes in the calorimeters, and evaluating the energy loss from the calorimeter measurements or from a parameterization as a function of η and ϕ. After a refit at the vertex the track parameters are used for further analysis.

When dealing with data already selected by the trigger the first two steps (`MooMakePhiSegments`) and (`MooMakeRZSegments`) can be substitute with *ad hoc* makers that seed the track search in the regions selected by the trigger, by using the RegionSelector. We refer to this way of executing as ``seeded mode''. More details regarding the implementation of the Moore package in HLT are given in [13-43].

13.4 Signatures, rates and efficiencies

In the following subsections, we will present a preliminary assessment of the physics performance using algorithms for LVL2 and EF applied to representative final-state classes: electrons and photons; muons; jets, taus and missing ET; b-jets; and B-physics. This broad classification stems from the physics goals of the ATLAS experiment, as explained in Chapter 4 and in [13-1]. Whenever possible, a lot of care has been paid to include the realistic use of data formats, as

they come from the Read Out System, and their associated converters (as described in previous section). This is deemed to be particularly important, as those steps might represent a sizeable fraction of the available latencies, and imply a particularly smart use of software tools. Steering control (as described in Chapter 9) has also been employed, highlighting the flexible boundary between LVL2 and EF. Selection schemes are then derived, which contain the signatures used to decide whether or not to reject events. In order to maximize the discovery potential, the selection schemes generally only use inclusive signatures. With the exception of B physics, reconstruction of exclusive decays is not required and no topological variable (e.g. the calculation of invariant masses from a combination of several high-p_T objects) is used in the selection, although this is technically feasible at LVL2 or in particular in the EF (e.g. to select $Z \rightarrow l^+l^-$ decays exclusively).

It is worthwhile underlying that system performance (*e.g.* execution time, amount of data needed) is one of the major requirement in the HLT selection, to comply with the constraints imposed by the on-line environment and resources. In this chapter some indication of the compliance with those requirements will be given for representative selections, whilst a more detailed analysis of the different contributions to those figures can be found in Chapter 14. Since obviously system and physics performances are correlated, in general all results have been achieved by trying and optimise them concurrently.

13.4.1 e/gamma

In the present view of the ATLAS trigger menus, the inclusive electron and photon triggers are expected to contribute an important fraction of the total high-p_T trigger rate. After the selection in LVL2 and the EF, the remaining rate will contain a significant contribution from signal events from Standard Model physics processes containing real isolated electrons or photons (b,c $\rightarrow$ eX, W $\rightarrow$ eν, Z $\rightarrow$ ee, direct photon production, etc.).

The electron and photon triggers can be viewed as a series of selection steps of increasing complexity. At LVL1 electrons and photons are selected based on the calorimeter information. After receiving the LVL1 electromagnetic (e.m.) trigger RoI positions, the LVL2 trigger performs a selection of isolated e.m. clusters using the full calorimeter granularity within the RoI which has the size $\Delta\eta \times \Delta\phi$=0.2 $\times$0.2. To calculate the cluster-E_T the latest calibration available is used (see Section [13.3.3.5]). Electrons are selected based on the cluster-E_T and shower-shape quantities that distinguish isolated e.m. objects from jets. A further, more refined calorimeter-based selection may classify the e.m. cluster as a LVL2 photon trigger object. In the next step, electrons are identified at LVL2 by associating the e.m. cluster with a track in the Inner Detector. Typically, track candidates are found by independent searches in the TRT and SCT/Pixel ('Precision') detectors in the region identified by the LVL1 RoI. Details of the different LVL2 tracking algorithms, which are currently being studied are described in Sections [13.3.3.1], [13.3.3.2], [13.3.3.4]. For the results presented in the next section, however, only one tracking algorithm, namely IDscan [13.3.3.1], has been considered and work is in progress to evaluate the others. For electron candidates a position and momentum matching between the track and the cluster is required. Electron candidates passing these selection criteria are retained to be examined by the EF.

At the EF electrons are selected by a very similar strategy as at LVL2 using the information in the calorimeters and the inner detector. For track reconstruction the results using xKalman are presented here and work is in progress to evaluate the alternative track reconstruction program iPatRec. First results show very similar behaviour between these two algorithms. At the EF, the

main differences with respect to LVL2 arise from the availability of improved calibrations and the possibility to use more sophisticated reconstruction algorithms with access to the full detector data. The improved performance results in sharper thresholds and better background rejection. In order to avoid a bias introduced by different reconstruction algorithms, the EF will select events using as far as possible the offline reconstruction algorithms. Though time consuming parts such as sophisticated bremsstrahlung recovery procedure will not be done at the EF level. Using the offline reconstruction in the EF implies these algorithms have to comply with the stricter EF requirements in terms of robustness and system performance. Currently this is not yet achieved. However, work is in progress to change the algorithms accordingly. The present study uses the currently available ATLAS offline reconstruction software as discussed in Section [13.3.4] as a prototype of the future EF code. The criteria to identify electrons and photons need to be softer at the EF level in order not to loose events prematurely. In previous studies [13-55] and [13-57], the offline electron and photon selection has been applied using the same identification criteria as the offline selection just leaving out few “critical” criteria. For example a track veto for non-converted photons has not been applied on the EF level because it requires a good control of the fake tracks in the inner detector and thus, a very good understanding of the tracking performance especially in the presence of pile-up. A more realistic EF electron selection has been used for the studies presented here.

In the following, the physics performance of the selection of electrons by the HLT will be reviewed. The electron and photon selection by the HLT has been studied very extensively in the past by the HLT and are reported in [13-55],[13-56],[13-57]. Here only the main performance issues are discussed obtained using the new software within the HLTSSW. A more complete overview can be found in [13-58]. Work is in progress to fully access the physics such as the photon selection, which will be reported on at a later stage. The system performance part of this selection is discussed briefly in Section 13.4.1.2 here but is as well documented in Chapter 14.

13.4.1.1 HLT Electron Selection Performance

The performance of the electron and photon triggers has been estimated for single electrons and photons, and for some standard physics channels (e.g. $Z \rightarrow ee$, $W \rightarrow e\nu$, $H \rightarrow 4e$). The physics performance has been characterized in terms of efficiency for the signal channel, and the rate expected for the selection. Here only the single electron trigger is considered. For the double object triggers the adequate datasets could not be reconstructed and analysed in time for the TDR. The rates and efficiencies have been obtained using fully simulated single electrons and around $2*10^6$ di-jet events at low and design luminosity (see Ref. [13-62]). A higher statistics of 10^7 di-jets is available, which still needs to be analysed. The events were generated using PYTHIA 6.203 [13-59]. It should be noted that compared to older studies the cross section for QCD di-jets with E_T>17GeV has increased by 50% [13-60]. However, a comparison with the fast simulation [13-60] has shown that at the same time the number of isolated electron has decreased by 50% due to a decrease in the number of electrons from b and c-decays. Compared to previous studies a more up to date detector geometry has been used. The main changes are the increased material distribution in the inner detector. The inner detector now contains more material especially due to the support of the first b-layer. Therefore, bremsstrahlung effects for electrons are becoming more important, especially in the end-caps. In addition, the correct magnetic field map is now used in the simulation instead of a constant field of 2T. This change affects the tracking performance in the end-caps. More details on the produced datasets can be found in [13-62].

In general, events with electrons and photons are selected on the basis of single high-p_T objects or of pairs of lower-p_T objects. The current physics performance of the single electron triggers is summarized here and documented in detail for the three trigger levels in Ref. [13-61] and Ref.

[13-58].

Table 13-1 Performance of the single electron HLT trigger at low luminosity. The results are presented in a single sequence. 'Matching' refers to position and energy–momentum matching between calorimeter clusters and reconstructed tracks (at LVL2 only precision tracks are used). The efficiencies are given for single electrons of p_T = 25 GeV a low luminosity over the full rapidity range $|\eta| < 2.5$. The efficiencies and rates are given with respect to a LVL1 output efficiency of 95% and a LVL1 rate for e.m. clusters of ~12 kHz. These numbers are very preliminary and will improve after some more studies.

Trigger Step	Rate [Hz]	Efficiency [%]
LVL2 Calo	xxx ± 24	92.1 ± 0.5
LVL2 Tracking	xxx ± 11	xx.4 ± 0.8
LVL2 Matching	86 ± 10	85.6± 0.4
EF Calo	29 ± 5	82.0 ± 0.4
EF ID	23 ± 5	73.2 ± 0.4
EF Matching	17 ± 4	70.1 ± 0.5

An overview of the current performance of the single isolated electron HLT algorithm is summarized in table [13-1] as a function of the main steps in the LVL2–EF trigger chain for the low luminosity scenario. The trigger steps have been factorized by detector in order to show the rejection that each stage contributes to the trigger. It should be noted, that there are strong correlations between the different selection criteria. The overall reduction in rate achieved by LVL2 is a factor of 140 for a loss of efficiency of 14% with respect to LVL1. The additional rate reduction provided by the EF amounts to a factor of 1.7 for an additional efficiency loss of ~20%. This shows that the HLT selection is very powerful. Compared to previous results (see Ref. [13-57]) the LVL2 selection is reducing the rate by a factor of 140 compared to **XXX** quoted in Ref [13-57]. Consequently at the EF level the additional rate reduction is very small. 2-3% of the events are rejected with the current selection, which would be accepted by the EF. These numbers still need to be understood better and further cross-checks are needed to insure the selection doesn't reject too many 'real' electrons prematurely. As a cross-check the electron efficiencies have been extracted using fully simulated $W \rightarrow e\nu$ events. With respect to LVL1, an efficiency of $(85.0 \pm 0.4)\%$ after LVL2 and $(72.7 \pm 0.5)\%$ has been obtained at low luminosity for W events for which the electron has a E_T>25GeV at generator level. These values are in agreement with the values given in table [13-1] for single electrons of E_T=25 GeV. The analysis of the performance of the single electron trigger e30i at high luminosity is in progress. First results using only the EF selection give a rate of 165 Hz for an efficiency of ~72% w.r.t. LVL1.

The rates quoted here are a factor of 2.5 lower than in the past for an electron efficiency which is 10% lower. The aim at the EF is to accept electrons with an overall efficiency of around 80% (corresponding to an efficiency of ~84% w.r.t. LVL1) in order not to cut too hard on physics. To reach a given rate of 40Hz at low luminosity, which is the rate obtained in the older studies, the electron efficiency would be ~83% w.r.t LVL1 as can be seen in the table extrapolating the results given after the EF calorimeter selection. From this it can be concluded that a final rate of 40Hz at low luminosity for a 80% electron efficiency as obtained in previous studies is still valid using the new selection algorithms, updated detector geometry and the changes at generator level. These results are preliminary and work is in progress to optimise and cross-check the current selection cuts. For example the cuts will be tuned in a more optimised way as a function of rapidity and additional criteria such as isolation in the inner detector around the electron track are

being studied. The statistices for the surviving events are quite low and the full statistics of the di-jet events still needs to be analysed.

13.4.1.2 HLT Strategy, Algorithm Optimisation and the LVL2–EF Boundary

The system performance has been evaluated for LVL2. The algorihms exectution time of both the cluster reconstruction as well as the track reconstruction using IDscan takes on average a 1-2 ms on a 2 GHz machine. For the calorimeters the data preparation step currently takes around 3 ms per RoI. In the moment the overal latency for the data preparation and algorithms is within 10ms on a 2 GHz machine and already now well within the latencies. The important and not expected outcome of these timing measurements is that most of the time is spent in the data preparation and not in the algorithmic part. This shows the importance to have the correct treatment of the data as seen by the trigger simulated correctly in the trigger software chain. For the EF no timings are given here, because the offline algorithms cannot yet run easily in the HLTSSW. However, first timing measurments have revealed bottlenecks in the reconstruction algorithms and work is in progress on this level. As for LVL2 it is planned to profile not only the algorithmic part, but as well the data access. For the LArg calorimeter first timing measurement have shown that on the EF the data unpacking takes ~0.18 s. In case the RoI concept is applied on the EF as an alternative approach, which was as well tested in earlier studies, the same time as found for LVL2 is expected. More information can be found in Chapter 14**.** Studies under way which will be documented in future TDAQ notes.

The use of system resources in the electron HLT can be minimized by exploiting the modularity of the trigger. By ordering the trigger steps in such a way that events are rejected as early as possible, both overall processing times and data transfers are reduced. Factorizing the trigger algorithm components also provides flexibility to move the rejection power from LVL2 to the EF or vice versa, to optimize the following: the performance of the implementation of the algorithm; the robustness of the selection with respect to the rate; the load implied at each level; etc. These issues have been extensively studied in the past and are reported in Ref. [13-3]. They are not repeated here, and more information can be found in Ref. [13-58].

In case the incoming trigger rate is too high the easiest handle to reduce the rates is either to raise the energy threshold of the trigger menu item or by stricter selection criteria. However, this implies an additional loss in efficiency for physics signals. This loss in physics can then partly be recovered by more selective triggers. There are long-term ongoing studies together with the physics community to assess the impact of such changes in order to be prepared for this case and be able to do the best to retain physics. The preferred and easiest way to reduce the rate is to raise the energy thresholds. The LVL1 rate is dominated by the contribution from single high-p_T e.m. objects. As an example, raising the thresholds by E_T=5 GeV of the single electron trigger would yield in a final HLT rate of **xxx (xxx)** at low (design) luminosity. This is also seen in Figure 13-6, which illustrates the impact of raising the threshold for the single-electron HLT selection at low luminosity (nominal threshold of 25 GeV) for $W \rightarrow e\nu$ events. As the threshold is increased, the efficiency to select these events decreases gradually. The impact on other physics signal such as $Z \rightarrow ee$, and $H \rightarrow 4e$ are discussed in Ref. [13-58]. As illustrated above, the proposed strategy contains considerable flexibility. Various possibilities exist to reduce the required computing resources or to improve the physics performance. For many channels of interest, the selection scheme also provides considerable redundancy. More details on the trigger selection strategy have been reported in Ref. [13-3] and some more details are given in Chapter 4.

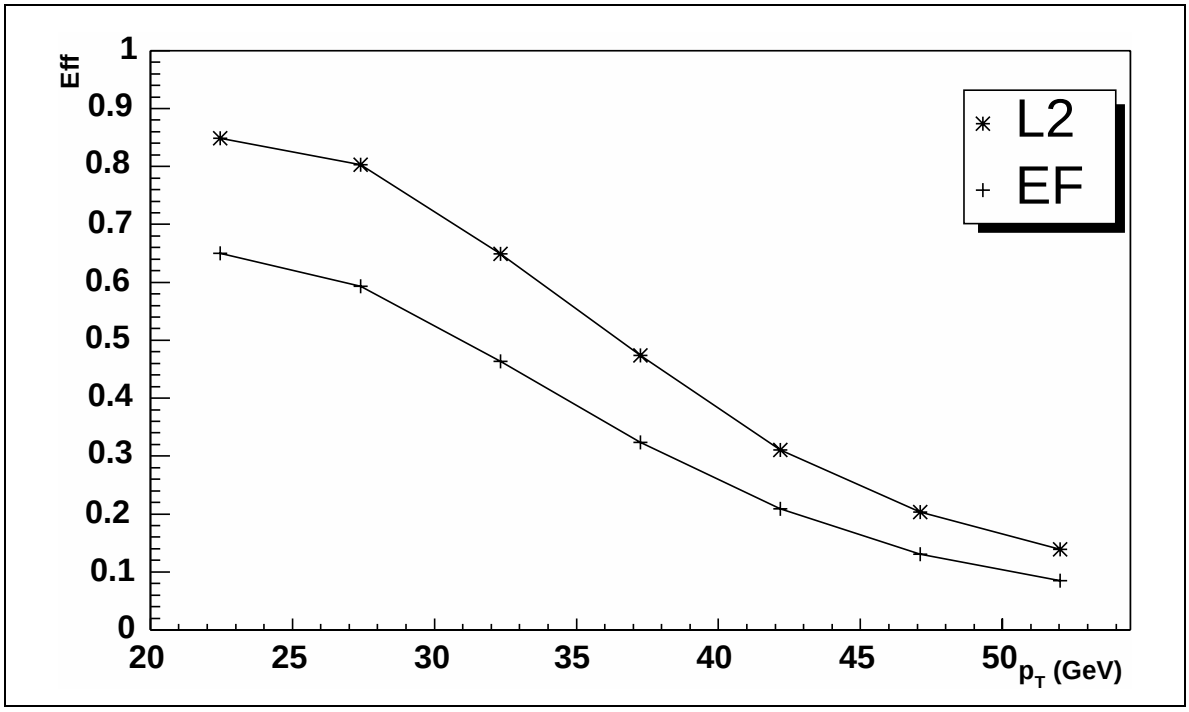

Figure 13-6 Efficiency to select $W \rightarrow e\nu$ events at LVL2 and EFas a function of E_T for the single electron trigger e25i at low luminosity.

13.4.2 Muon selection

The main purpose of the high-level muon trigger is the accurate reconstruction of muon tracks in the RoIs indicated by the LVL1 muon trigger. LVL2 and EF must reject low-p_T muons, secondary muons produced in the in flight decays of charged pions and kaons, and fake muons originating from the cavern background. The EF must be able to reconstruct additional muons present in the event not reconstructed or selected by the LVL2 trigger.

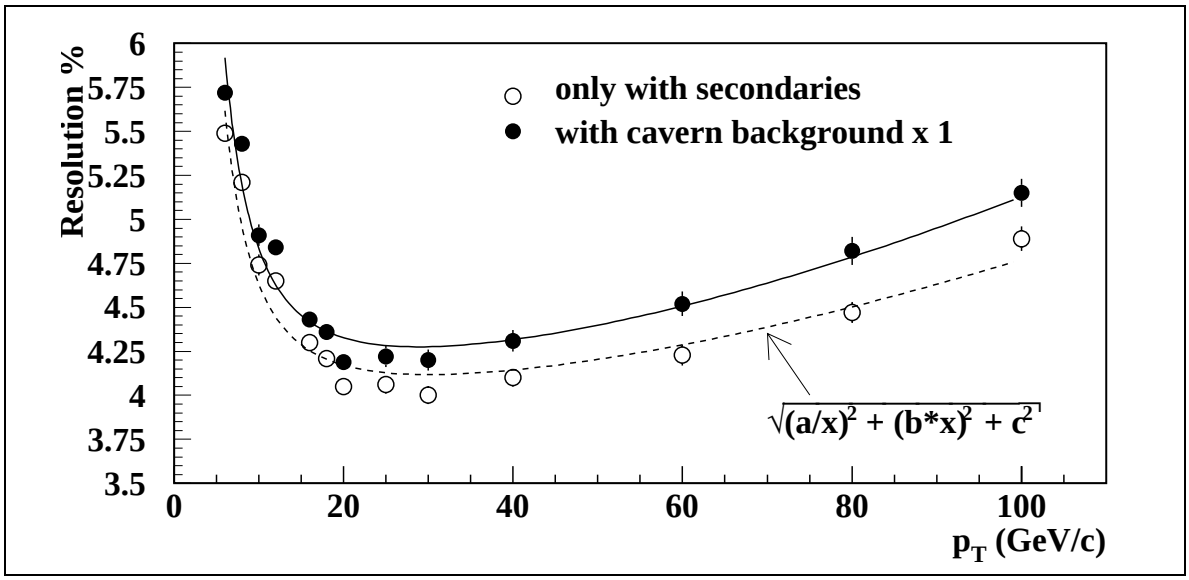

Figure 13-7 The p_T resolution of the muFast algorithm as a function of muon p_T.

Whilst the LVL1 trigger system uses only hits from the dedicated trigger detectors (RPCs in the barrel and TGCs in the endcap), the LVL2 and EF has access to the full measurements of the Muon Spectrometer, including in particular the data from the Monitored Drift Tubes (MDTs). This allows the best muon track reconstruction. The high background environment in the Muon Spectrometer demands algorithms with robust and fast pattern recognition capable of rejecting hist induced by the cavern background.

The tracks found in the LVL2 Muon Trigger are extrapolated for combination with the Inner Detector and the Calorimeter. Matching between muon tracks measured independently in the Muon System and the Inner Detector selects prompt muons and reject fake and secondary muons. This is important in particular for the B-physics trigger in low-luminosity running, for which the selection of prompt low-p_T muons events defines the input of the B-physics trigger algorithm. The studies presented in this section are limited to the barrel region ($|\eta|<1$).

13.4.2.1 The Physics Performances of LVL2 Muon algorithms

The physics performances of the LVL2 muon trigger algorithms have been presented and discussed in detail in the HLT/DAQ/DCS Technical Proposal [13-3].. The algorithm muFast has been implemented with no important changes in the code used to reconstruct the muon tracks. The main difference of the present version of the code with respect the one studied in the TP is the use of the "LVL1 Emulation" to identify among the RPC hits the ones used by the LVL1 trigger to select muon candidates. We expect this to have very little impact on the overall muon reconstruction efficiencies at LVL2, and no effect at all on the muon transverse momentum resolution. Here we recall the most relevant results obtained [13-34]. The p_T resolution of reconstructed muons is crucial to the selection efficiency and to the rejection of low p_T tracks that can be achieved at LVL2. The p_T resolution of the muFast algorithm is shown as a function of p_T in Figure 13-7. As shown in the figure, the resolution ranges between 4.0% and 5.5% for muon in the p_T interval 6-20 GeV/c. These results are well compared with the transverse momentum resolution obtained by the offline muon reconstruction program MUONBOX [13-64].

The total rates after this algorithm, including the rejection provided by the LVL1 selection, have been evaluated by convolving the algorithm efficiency as a function of p_T with the muon differential cross section production of the dominant physics processes. Where the available statistics are too low (in particular for the high-p_T rate calculation) to evaluate the efficiency, the lowest p_T at which an efficiency estimate has ben possible (p_T=10 GeV/c) is assumed conservatively to constitute a plateau extending down to the lower limit of the p_T acceptance (p_T=3 GeV/c in the barrel). The rates from π/K decays are calculated using the predicted cross-sections from the DMPJET program, and would be lower by about 50% if the PYTHIA prediction were used. The total rates after LVL2 are shown in Table 13-2[1]. Preliminary studies of the trigger rate arising from the cavern background as predicted by the FLUKA package have been done. The probability that a fake LVL1 muon trigger is accepted by the LVL2 is below 10^{-2}. This upper limit is sufficient to neglect the contribution from fake muons.

Table 13-2 Total output rates [kHz] of the standalone LVL2 muon trigger after application of the muFast algorithm for the 6 GeV/c low-p_T threshold at 1×10^{33} cm^{-2} s^{-1} and 20 GeV threshold at the design luminosity.

Physics Process	low-p_T	high-p_T
π/K decays	3.00	0.07
b decays	0.90	0.09
c decays	0.50	0.04
W$\rightarrow\mu\nu$	0.003	0.03
cavern background	negligible	negligible
Total	**4.40**	**0.23**

The new implementation of muFast in the trigger framework has been used preliminarily to reconstruct muons with p_T of 20 GeV/c. Figure 13-8 shows the p_T resolution $(1/p_T\text{-}1/p_T^{true})/(1/p_T^{true})$ obtained by muFast with the implementation in the new PESA software. The value found is 4.0 and it is identical to the result obtained with muFast in the old implementation. This result supports the expectation of no relevant physics performance changes with respect to past results [13-34].

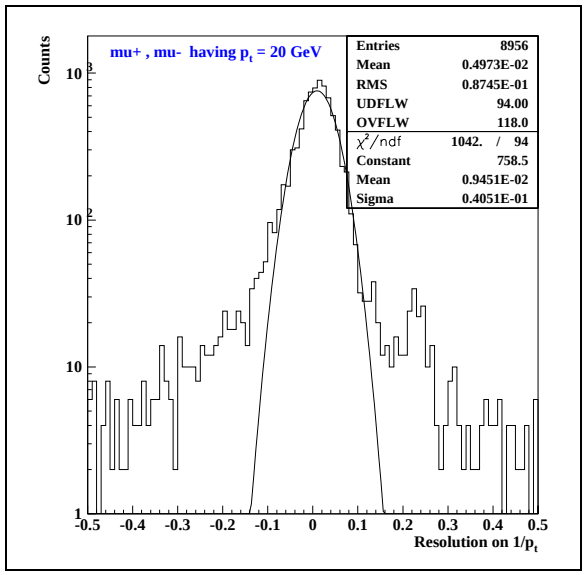

Figure 13-8 Transverse momentum resolution $(1/p_T\text{-}1/p_T^{true})/(1/p_T^{true})$ of 20 GeV muons reconstructed with mufast in the new PESA framework (no cavern background). The tails arise from events with large Coulomb scattering or from poor momentum reconstruction indiced by delta-rays emitted by the muon.

The implementation of the muComb algorithm in the new framework is still on-going. Also in this case the algorithm under implementation is the same discussed in the HLT/DAQ/DCS TP, and therefore we can recall here the results already found. Figure 13-9 shows the combined reconstruction efficiency of prompt and secondary muons from π/K in flight decays, as a function of the muon p_T, where the standalone codes from muFast and the LVL2 Precision algorithm have been used. The requirement of a good muon track matching (z/ϕ and p_T matching) reduces the low p_T trigger rate to 1.0 kHz: a factor three reduction compared to the rate from the muFast algorithm. Including the further reduction in rate due to the increase in p_T resolution for prompt muons, the total rate from the muComb algorithm is 2.1 kHz from muons with p_T>6 GeV/c..

1. $W\rightarrow\mu\nu$ cross section found with PYTHIA for p_T^{μ}>3 GeV and $|\eta_\mu|$<2.7 : 9.56 pb.This corresponds to 100 Hz at 1×10^{34} cm^{-2} s^{-1}

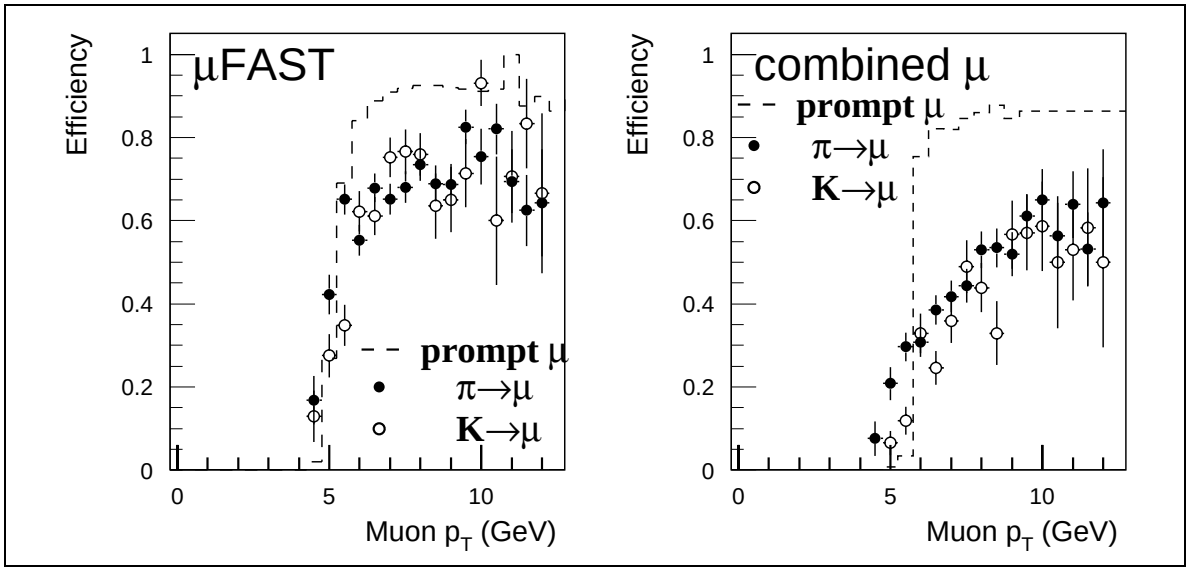

Figure 13-9 The efficiency with respect to LVL1 of the combined reconstruction at LVL2 for prompt muons and for muons from π/k decays. The left hand plot shows the efficiency of the standalone algorithm muFast and the right hand plot shows the efficiency of the combined muon algorithm muComb.

13.4.2.2 The Physics Performances of the Muon Event Filter

The physics performance of the MOORE EF package has been evaluated with simulated single muon samples, with no cavern background, in the p_T range 3 to 1000 GeV/c. Here we have considered the *fully wrapped* version. This is fully equivalent, from the reconstruction point of view, to the offline version. Figure 13-10 shows the muon track reconstruction efficiency as a function of generated p_T in the barrel region. The bullets show the efficiency relative to the reconstruction in the standalone muon spectrometer ("Moore"), while the square symbols show the efficiency of the reconstruction when the track is extrapolated to the interaction region ("Muid"). The energy loss in the calorimeter has been parameterized as a function of η and ϕ. The loss of efficiency of "Muid" at low p_T is due to the failure of the extrapolation of muons that exhibit in the spectrometer a transverse momentum of a few GeV/c. Muon with p_T larger than 8 GeV/c are reconstructed with efficiencies above 95%, and this is equivalent to the result showed by MUONBOX in the Physics TDR [13-1],[13-64].

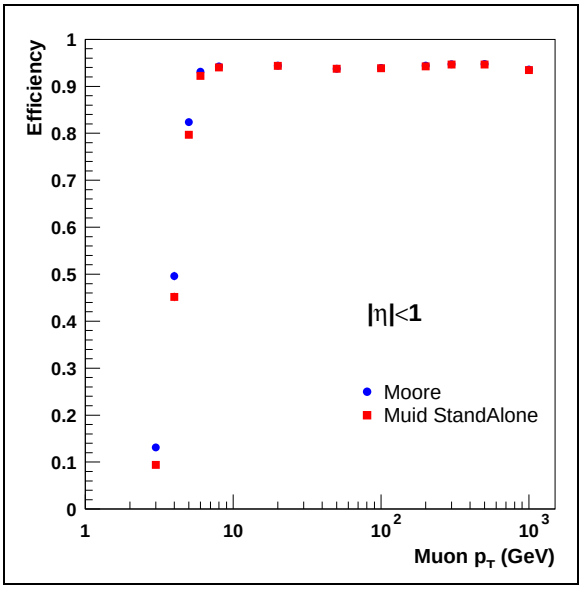

Figure 13-10 Single muon reconstruction efficiency as a function of the generated pT in the barrel region; bullets: standalone reconstruction (Moore); squares: reconstruction at the nominal interaction vertex point.

Figure 13-11 shows the muon transverse momentum resolution of the MOORE package again in the standalone muon spectrometer and when the track is extrapolated to the interaction vertex. Again, the result are coherent with the result found in past studies. Concerning the physics per-

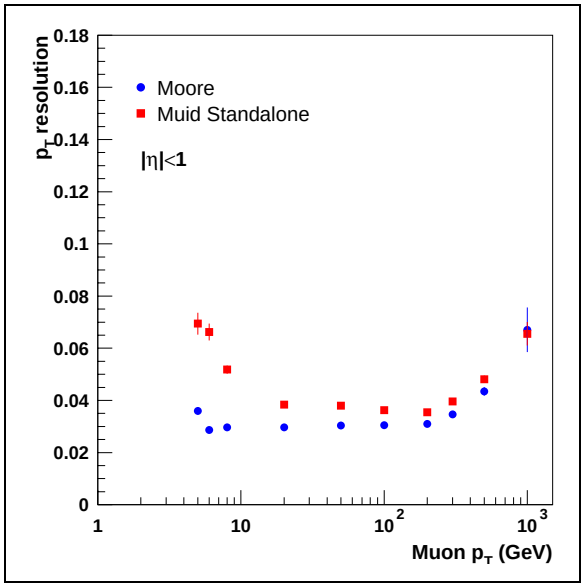

Figure 13-11 *Fully wrapped* version: transverse momentum resolution as a function of the generated p_T.

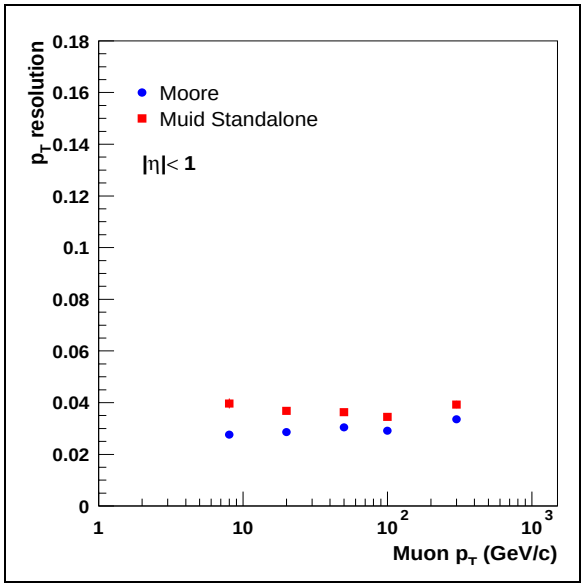

Figure 13-12 *Seeded* version: transverse momentum resolution as a function of the generated p_T.

formances of the *seeded* [13-43] version, we don't expect significant differences with respect the full version, since the reconstruction and fitting methods are the same of the ones used by the offline package. This is supported by Figure 13-12 where the transverse momentum resolution of the seeded version is shown for p_T=20 and 300 GeV/c. A more complete presentation of the MOORE physics performance is reported in [13-42].

13.4.2.3 The Timing Performances of the Muon Algorithms

The muFast trigger algorithm has been benchmarked on 2 GHz machine. The event sample consisted of about 430 events of single muons in the barrel region with 100 GeV transverse momentum, with and without cavern background. As shown in Figure 13-13 the average processing time of muFast is about 1 ms with a r.m.s. of 0.2 ms and with all events within 4.1 ms. Figure 13-14 shows the correlation of the processing time with the RPC activity; as it can be seen a linear type relation. Finally, the composition of this overall timing, is shown in Figure 13-15; as it can be seen, the processing time is dominated by the Level-1 Emulation (about 0.8 ms). In presence of the cavern background (luminosity $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and nominal intensity, i.e. safety factor=1) the muFast timing increase by about 100 microseconds.

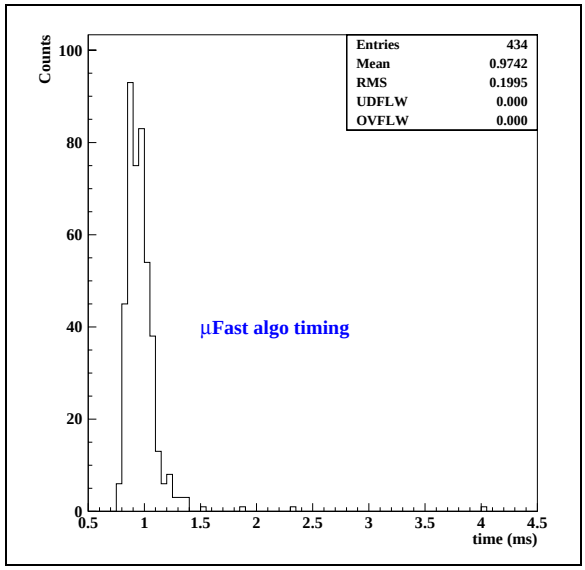

Figure 13-13 Processing time of muFast.

The timing performance of the Moore algorithm, both for seeded and wrapped mode, have also been evaluated on a 2 GHz machine. The code has been built in optimized more. The event sam-

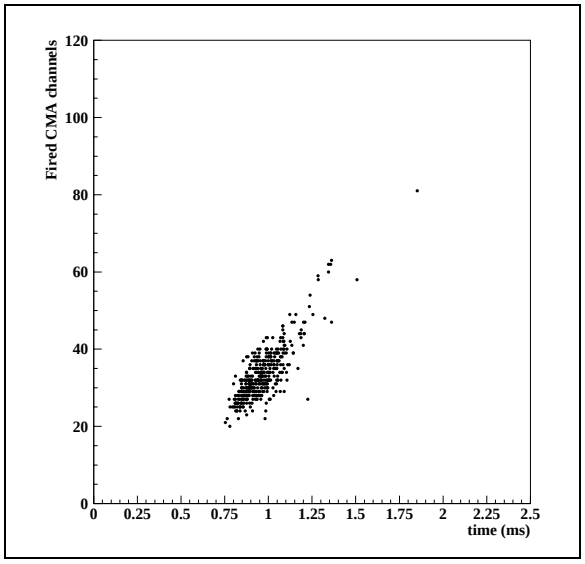

Figure 13-14 Correlation between the muFast processing time and the number of the fired channels in the RPC readout.

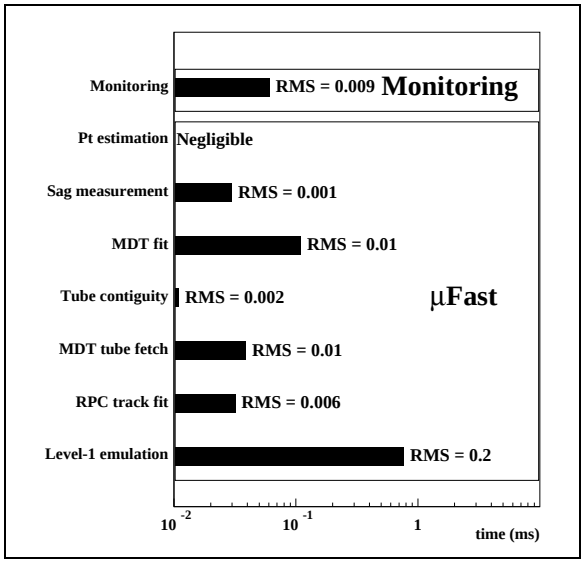

Figure 13-15 Contributions to the muFast processing time. The bar length shows the average processing time; the RMS is given in ms.

ples consist of around 500 events each.In table {vedi SECONDA tabella} the average execution times per event both for seeded and wrapped mode have been included. The timing includes the procedure of extrapolation to the vertex of the reconstructed track. Moreover, until now a conservative approach has been adopted. This means that the timings include also the accesses to the data, the data preparation, and the region selector accesses.

Table 13-3 Overview of timing tests (n. d. r. average on one second window)

Sample	Time (seeded mode) msec		Time (wrapped mode barrel) msec	
	< >	RMS	< >	RMS
8 Gev	86	69	80	63
20 Gev	65	36	66	43
50 GeV	68	68	68	51
100 GeV	71	60	76	69
300	75	60	81	88
100 GeV + X1	775		2.8 sec	
100 Gev + X5			24 sec	9 sec

The general behaviour of the average execution time for the different samples is represented in Figure 13-16 for both seeded and wrapped mode. In order to show the impact of the extrapolation to the vertex we have plotted the execution times for the track reconstruction inside the MuonSpectrometer (Moore), the execution times for extrapolating the track to the vertex (MuID), and the sum of the two (total). The execution times are rather flat over the analysed p_T range.The latency time for the whole reconstruction, data access and data preparation of single muon is in average below 100 msec. For those events we do not expect a large difference running in wrapped and seeded mode, since the data set that is accessed is quite similar.

The plot has been performed averaging only on events for each sample that register execution

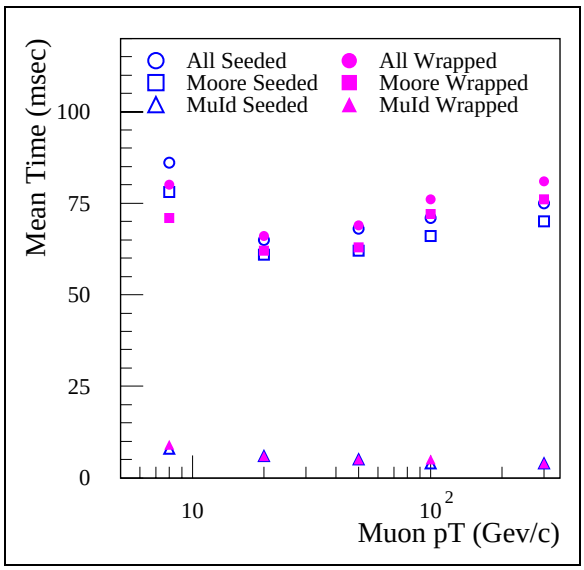

Figure 13-16 Average execution time (in msec) for different p_T values (in GeV), obtained with Moore, with MuId standalone and with both of them, in seeded and wrapped mode.

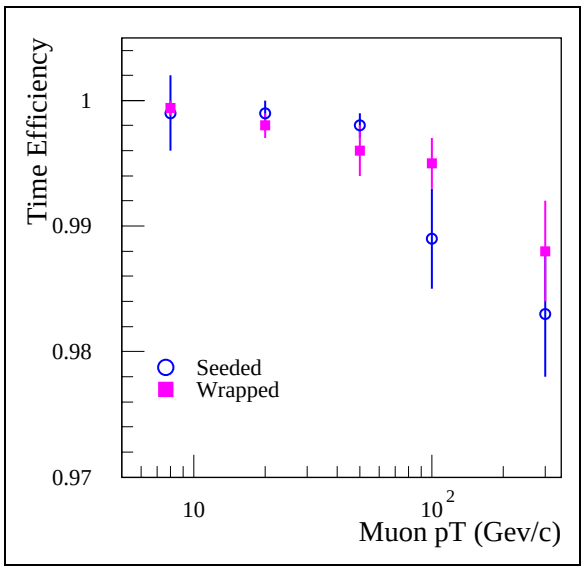

Figure 13-17 Time efficiency for seeded and wrapped mode for different p_T values.

times below one second. In order to show the impact of events that register longer execution times we define a time efficiency as the ratio between the number of reconstructed tracks in one second and the number of Region of Interest. The plot is shown in Figure 13-17 for both seeded and wrapped mode.

13.4.3 Tau/Jets/$E_T$miss selection

13.4.3.1 The Tau Trigger

A major Standard Model source of tau leptons in ATLAS will arise from *W/Z* decay sources: $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$. The tau lepton will also play a key role in the search for new physics. In the MSSM the heavy scalar (H) and pseudoscalar (A) Higgs boson decays to a tau-pair are strongly enhanced with respect to Standard Model Higgs boson case. Also, a key decay channel for the charged Higgs boson is $H^{\pm} \rightarrow \tau\nu$.

The identification of the hadronic decays of Tau leptons is based on the selection of narrow isolated jets with low multiplicity in the tracking system. The shower isolation and shape are calculated for both the e.m. and hadronic calorimeters separately. The fraction of energy deposited by the tau-jet in the e.m. calorimeter has a mean value around 60%. The hadronic shower is broader in the hadronic calorimeter than in the e.m. calorimeter. Thus the jet information obtained from the e.m. calorimeter is more selective than that from the hadronic calorimeter. At LVL1 the tau trigger described above would have similar inputs and much of the same logic, as the electron / photon trigger. A detailed [13-44] description of the tau trigger studies presented below is given elsewhere.

13.4.3.1.1 The First Level Tau Trigger

The motivation for a LVL1 tau calorimeter trigger is manyfold, both in conjunction with electron, muon or missing E_T signatures to enhance Z, A or W covergae, and for calibration purposes. Narrow tau jets containing 1(3) tracks give rise to narrow isolated energy depositions in the calorimeters. It is envisaged that an isolation requirement will be a valuable part of the tau trigger at all levels.

The e/gamma/tau LVL1 algorithms are described in detail elsewhere [13-2][13-3]. The LVL1 Tau/hadron calorimeter trigger is based on a 4 x 4 array of "trigger towers" in the electromagnetic and hadronic calorimeters (within the region $|\eta| < 2.5$) where the tower granularity is (0.1 x 0.1) $\Delta\eta$ x $\Delta\phi$. A core E_T is defined in the trigger algorithm as the sum of the electromagnetic and hadronic E_T in a 2 x 2 trigger tower. The trigger algorithm is based on four elements: the trigger cluster(s), an e.m. isolation region, a hadronic isolation region and an "RoI cluster".

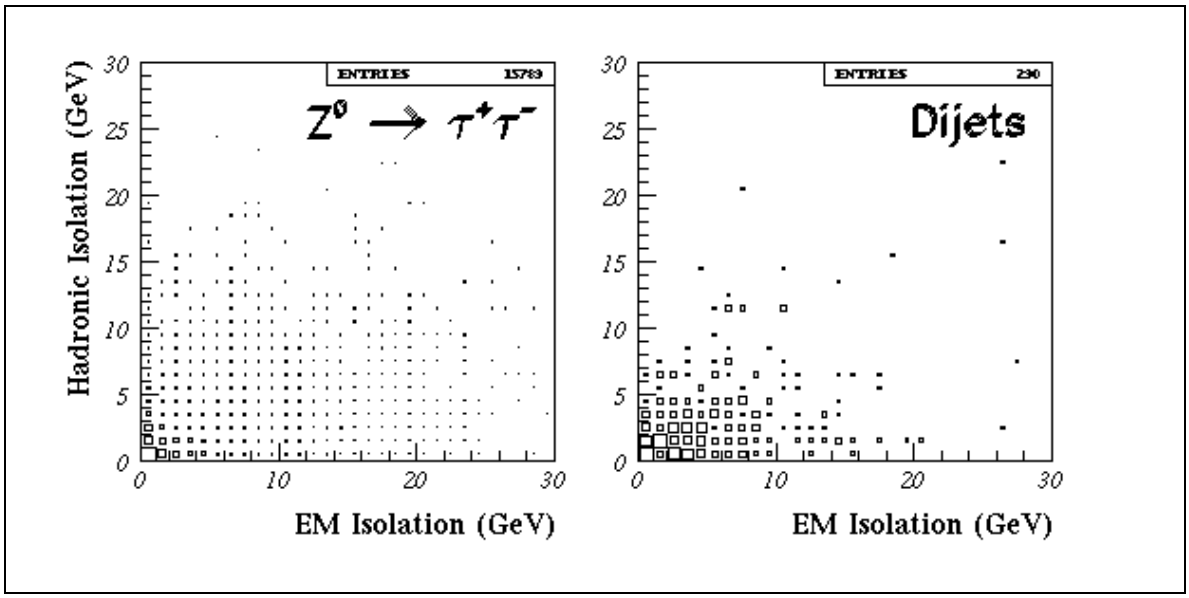

Figure 13-18 The hadronic isolation E_T vs. the e.m. isolation E_T (12 tower) for tau and QCD di-jets using the LVL1 trigger algorithm.

The relative power of the two isolation requirements, for the LVL1 tau/hadron trigger algorithm(2x1 e.m.+ 2x2 hadronic) is shown in the scatter plot Figure 13-18 created using $Z \rightarrow \tau^+\tau^-$ events and QCD di-jet events. The EM isolation is more powerful that the hadronic. In both plots a LVL1 algorithm has been employed with a core-E_T threshold of 20 GeV. The correlation between the isolation E_T values indicates that the hadronic isolation is of limited use in rejecting QCD jets. For example, the standalone tau trigger rate from QCD di-jets, at a luminosity of $1 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and neglecting the effect of pile-up, for a core-E_T of 20 GeV, a jet threshold of 20 GeV and an isolation cut at 10 GeV would be about 19 KHz.

13.4.3.1.2 The High Level Tau Trigger

The tau/electron trigger utilizes the e/gamma slice tools described in Section 13.4.1. The signal selection is tuned using events of the type $Z^0 \rightarrow \tau^+\tau^-$. Background evaluation is performed using fully simulated di-jet events. The LVL2 studies involve the verification of the LVL1 decision and, subsequently, tau identification using parameters that describe the shower shape defined

in layer one and two of the e.m. calorimeters. The LVL2 variables used in a previous analysis [13-45][13-48] could not be used due to changes in the e/gamma software implementation. Additional rejection of background jets can be achieved by using the information from tracks associated to the tau RoI.

The LVL2 algorithm is applied to LVL1 tau RoIs. Loose LVL1 cuts are chosen for the study presented here: a cluster E_T in excess of 20 GeV is required, and the e.m. and hadronic isolation thresholds are set to 10 GeV. LVL2 jet calibration is applied to the cells within the LVL1 RoI window. The energy weighted position of the tau-jet candidate ($\Delta\eta_\tau$ x $\Delta\phi_\tau$) is computed from all calorimeter cells within the LVL1 window. The first part of the LVL2 algorithm is the confirmation of the LVL1 decision. In order to do this the LVL1 algorithm described above is executed except that fine grained cell information is utilized.

As a first step, a check was made to ensure that LVL1 RoI coordinates are good approximations of the LVL2 tau coordinates, by measuring the distance between LVL1 RoI and the associated LVL2 cluster. After this successful check, for the LVL2 clusters associated with LVL1 RoIs the next step is to look at the LVL2 e/gamma calorimetric variables that have some power to select taus over QCD jets. Three variables were identified. The performance of the algorithm as a function of the shower shape variable R_{73} was examined first. This variable is defined as the ratio of the E_T contained in a 3x7 cell cluster to the E_T in a 7x7 cell cluster centred on the same seed cell, calculated for the 2nd layer of the e.m. calorimeter The distribution of R_{73} calculated for $Z^o \rightarrow \tau^+\tau^-$ and QCD di-jets is plotted in Figure 13-19.

The second variable studied, $F_{12,}$ is defined in the first layer of the e.m. calorimeter and functions as an isolation variable: $F_{12} = (E_{max1} - E_{max2})/(E_{max1} + E_{max2})$ where E_{max1} and E_{max2} are the maximum and next to maximum energy strips, respectively, in the first layer in the e.m. calorimeter: this quantity is defined for each LVL2 cluster. Is should be noted that the variables R_{37} and F_{12} are highly correlated. The third variable was, $F_{e.m.}$, the e.m fraction of the total energy.

Table 13-4 Tau and QCD jet trigger efficiencies for different R_{37} cuts and a defined LVL1 trigger condition. The right hand column shows the corresponding stand-alone tau trigger rate based on QCD di-jets with a Pythia generation threshold E_T of 35 GeV.

Cut	Tau Efficiency % (wrt LVL1 accepts)	QCD jet efficiency% (wrt LVL1 accepts)	LVL1/2 tau trigger rate using calorimetry (KHz)
LVL1 20-10-10			
$R_{37} > 0.75$	88.± 1	71 ± 1	11.1 ± 0.2
$R_{37} > 0.80$	83± 1	59 ± 1	9.4 ± 0.2
$R_{37} > 0.85$	75 ± 1	44 ± 1	7.0 ± 0.2
$R_{37} > 0.90$	62 ± 1	26 ± 1	2.3 ± 0.2

The variable with the most power to select taus and reject QCD jets was seen to be R_{37}. Since the F_{12} variable is highly correlated to R_{37}, it was not found possible to obtain useful further improvements in tau efficiency and jet rejection using F_{12} whilst maintaining reasonably high tau efficiencies. Likewise, it was not possible to employ $F_{e.m.}$ to gain further QCD jet rejection whilst maintaining a good tau efficiency.

Three cuts in the R_{37} variable were studied: $R_{37} > 0.75, 0.80, 0.85$. These cuts were chosen to maintain a tau efficiency greater than 75%. The efficiency for triggering on $Z^o \rightarrow \tau^+\tau^-$ and QCD di-jets, along with the corresponding stand alone tau trigger rate at LVL2, estimated using QCD di-jets, for the three R_{37} cuts mentioned are given in Table 13-4. The chosen set of LVL1 conditions, denoted 20-10-10, are the loose LVL1 cuts defined for the LVL1 study: in this case the core E_T is 20 GeV and the e.m. and hadronic isolation thresholds are both 10 GeV. This can modified by strenghtening isolation requirements or core E_T values (*e.g.*, to favour tau trigger rate analyses based on higher mass $A \rightarrow \tau^+\tau^-$). It should be noted that tau efficiencies reported here are not directly comparable to the electron ones of previous section, because they derive from an analysis of a Z sample.

Additional rejection of background QCD jets can be achieved by using the information from tracks associated with the tau LVL1 RoI. The track information at LVL2 was used to associate Inner Detector tracks, found using the "IDscan" algorithm, with the tau RoI by requiring that ΔR between the track and the LVL2 tau cluster direction, associated with the tau LVL1 RoI, obeyed the relation, $\Delta R \leq 0.3$. The inner detector track multiplicity distributions obtained for $Z \rightarrow \tau^+\tau^-$ and QCD di-jets are shown in Figure 13-20. The resulting LVL2 tau and QCD di-jet efficiencies, for a few useful combinations of calorimeter and track based cuts, are shown in Table 13-5.

Table 13-5 This table lists the tau and QCD jet efficiencies for various useful combinations of R_{37} and LVL2 track multiplicity cuts, for a specified LVL1 trigger condition. The right hand column shows the corresponding stand-alone tau trigger rate based on QCD di-jets with a Pythia generation threshold E_T of 35 GeV.

Cuts	Tau efficiency % (wrt LVL1 accepts)	QCD di-jet efficiency % (wrt LVL1 accepts)	Stand alone tau/hadron trigger rate (KHz)
LVL1: 20-10-10			15.6 ± 0.2
R_{37}>0.8, #LVL2 tracks<5	74 ± 1	34 ± 1	5.3 ± 0.2
R_{37}>0.8, #LVL2 tracks<4	70 ± 1	30 ± 1	4.7 ± 0.2
R_{37}>0.9, #LVL2 tracks<5	57 ± 1	18 ± 1	2.8 ± 0.2
R_{37}>0.9, #LVL2 tracks<4	54 ± 1	15 ± 1	2.3 ± 0.2

13.4.3.1.3 Tau Selection in the Event Filter

At the Event Filter stage access to the complete, calibrated, event is possible for the first time. In addition, the tracking information has been further refined to reduce spurious track segments and minimize tracking inefficiency. Thus, it is possible to refine the LVL2 decision. Existing offline studies of tau / hadron identification and jet rejection [13-49] provide the basis for the event filter (EF) trigger decision. Typical trigger criteria for tau / hadron jets with $E_T > 30$GeV and $|\eta| < 2.5$ are as follows:

- The jet radius computed using only the e.m. cells contained in the jet, R_{em}, must obey the inequality: $R_{em} < 0.07$;
- The difference between the E_T contained in cones of size $\Delta R = 0.2$ and 0.1 normalized to the total jet E_T, ΔE_T must obey the inequality: $\Delta E_T < 0.1$ (isolation fraction);
- The number of reconstructed charge tracks N_{tr} is equal to 1 or 3, pointing to the cluster (within a ΔR of 0.3)

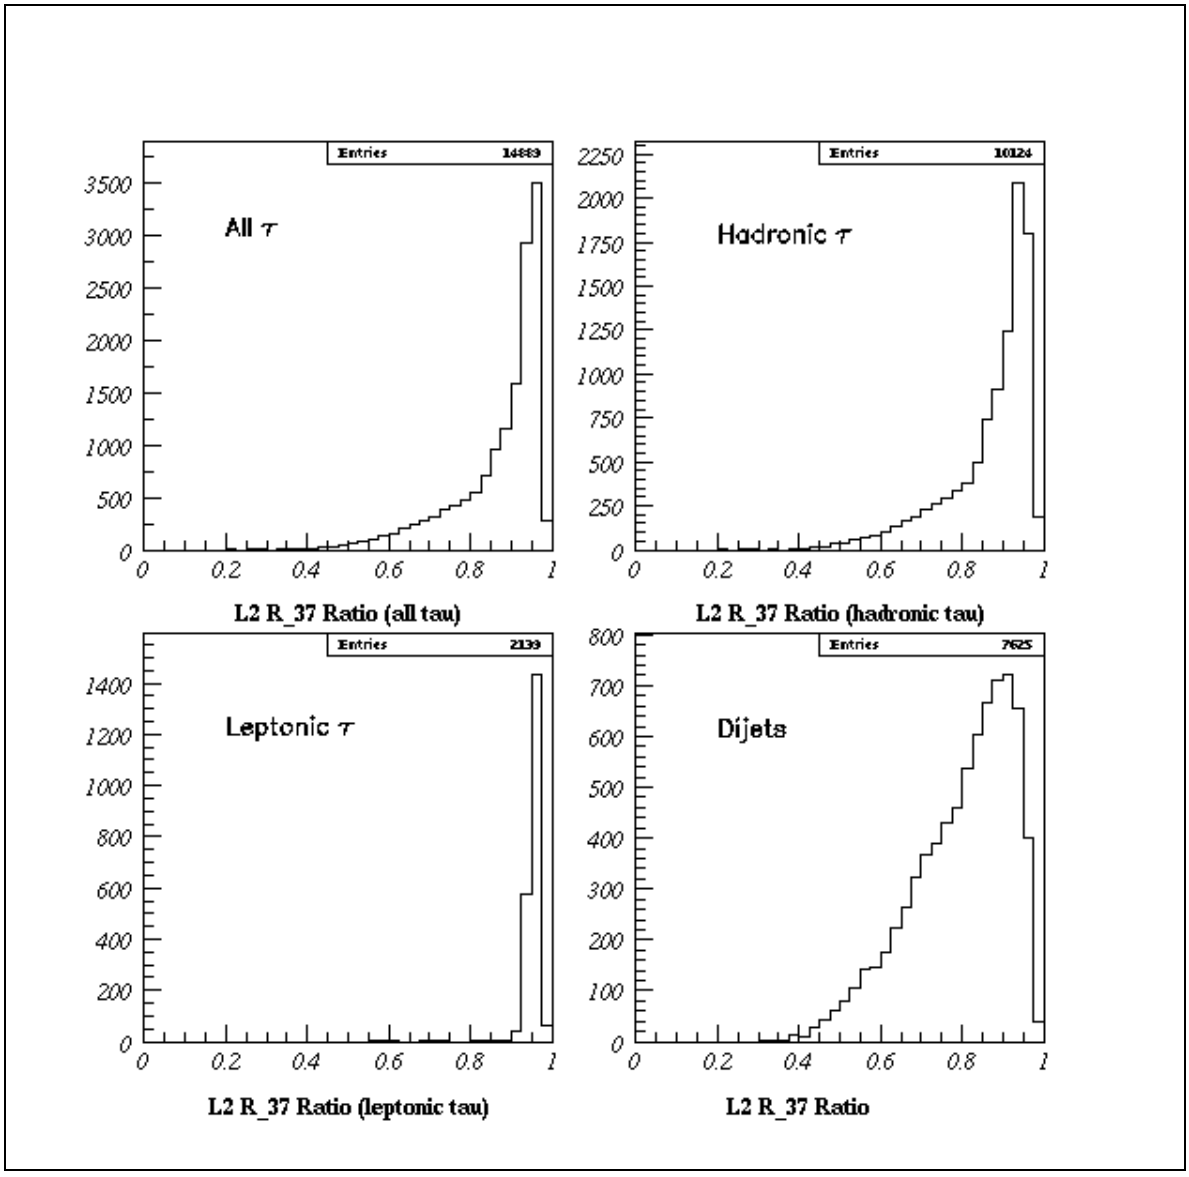

Figure 13-19 The ratio R_{37} for taus, from the process $Z^0 \rightarrow \tau^+\tau^-$ and for QCD di-jets.

Other tau identification variables are the tau hadronic E_T and a tau likelihood variable included in the "Taurec" reconstruction package in the ATHENA framework. The likelihood function utilizes simple analytical fits to the distributions of the above variables, as well as the E_T of the highest E_T track. The variables mentioned above still have power at the EF level, to reject QCD jets whilst retaining an adequate efficiency for selecting taus. An analysis of the rejection of QCD jets at the EF level is currently under way [13-44].

13.4.3.1.4 Jet Rejection at LVL2 Following a Tau +E_T-miss trigger

A method of improving the signal acceptance for final states involving taus, as well as retaining an acceptable trigger rate, is to combine the stand alone tau trigger with an E_T-miss trigger. The effect of a LVL1 tau +E_T-miss trigger on stand-alone tau HLT rates is given in Table 13-6 for one set of LVL2 trigger criteria and four trigger configurations: τ20 + xE25; τ20 + xE35; τ30 + xE25; and, τ30 + xE35. As can be seen from the table the LVL1 tau + E_T-miss trigger substantially en-

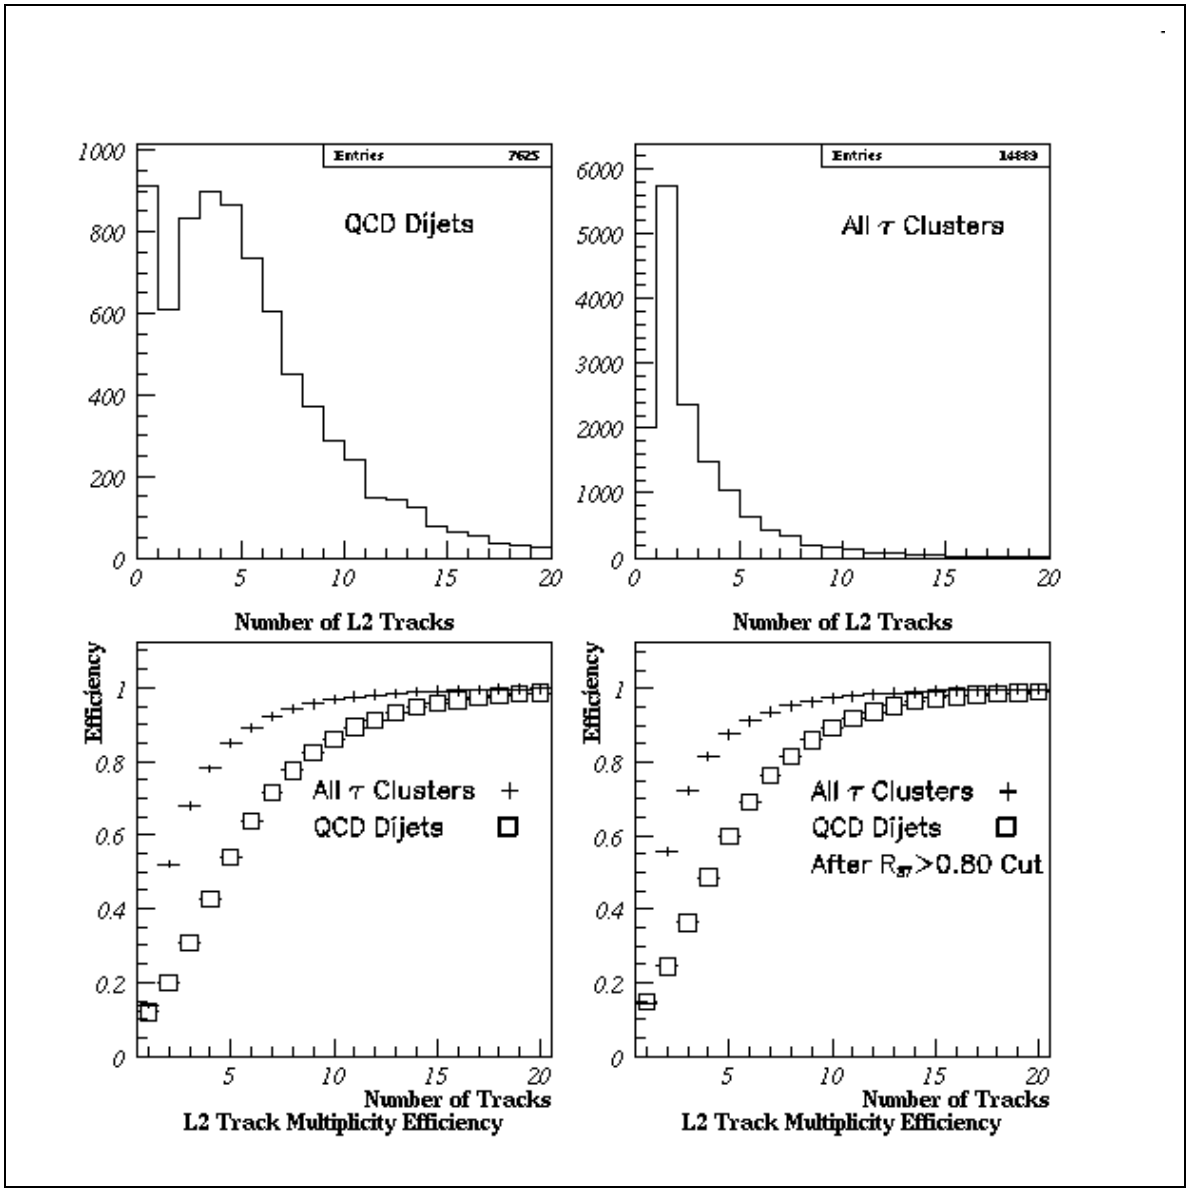

Figure 13-20 The top plots show the inner detector track multiplicities determined at LVL2 for $Z^0 \rightarrow \tau^+\tau^-$ &QCD di-jets The bottom plots show the variation of tau & jet finding efficiency with upper cut on track multiplicity.

hances the jet rejection and thus reduces considerably the stand alone tau trigger rate coming out of LVL2.

13.4.3.1.5 Preliminary estimates of Algorithm Timings

In order to obtain estimates of the trigger algorithm timings, the “Trigger Timer Service”, available from the standard HLT selection chain, was utilized. The LVL2 tau trigger used two basic triggering elements, the calorimetry and the tracking. In both cases the per event timings are found to be 25 ms for a 2 GHz CPU.

Table 13-6 Tau, and QCD di-jet efficiencies for various LVL1 t& LVL2 trigger criteria are listed in this table. In the right hand column the corresponding tau stand alone LVL2 trigger rates are given.

LVL1 Tau + E_T-miss trigger (GeV)	LVL2 Cuts	LVL2 tau eff. wrt LVL1 (%)	LVL2 jet eff. wrt LVL1 (%)	LVL2 stand alone tau trigger rate (KHz)
τ(20-10-10) + xE25	-	29 ± 1	10 ± 1	1.56 ± 0.16
τ(20-10-10) + xE35	-	11 ± 1	1.9 ± 0.2	0.30 ± 0.03
τ(20-10-10) + xE25	R_{37}>0.9 #LVL2tracks<5	15 ± 1	1.6 ± 0.2	0.25 ± 0.03
τ(20-10-10) + xE35	R_{37}>0.9 #LVL2tracks<5	5.1 ± 0.2	0.5 ± 0.1	0.07 ± 0.01
τ(30-10-10) + xE0	-	-	-	4.7 ± 0.1
τ(30-10-10) + xE35	-	20 ± 1	4.3 ± 0.4	0.20 ± 0.02
τ(30-10-10) + xE25	R_{37}>0.9 #LVL2tracks<5	23 ± 1	3.6 ± 0.4	0.17 ± 0.02
τ(30-10-10) + xE35	R_{37}>0.9 #LVL2tracks<5	9.6 ± 0.5	1.2 ± 0.2	0.06 ± 0.02

13.4.3.2 E_T-miss Trigger

Missing transverse energy will provide a distinct and important signature for new physics at the LHC. A precise and reliable measurement of E_T-miss requires good calorimeter performance and energy resolution, good linearity of response and hermetic coverage.

One of the basic building blocks of the LVL1 calorimeter trigger is the E_T sum in the calorimetry. The total scalar E_T, as well as its components, are computed in the Jet/Energy sum processor of the calorimeter trigger. An E_T-miss trigger is not implemented in the basic LVL1 inclusive triggers. However, it is an important part of the trigger when placed in combination with the tau/hadron, electron/photon and single jet triggers. A high level E_T-miss trigger is applied at the Event Filter level only, where access to the complete event, after calibration, is available. At this stage one can combine tau and E_T-miss triggers as one would in an offline analysis program, for example in the search for MSSM Higgs production in the channel $A/H \rightarrow \tau\tau$

The E_T-miss + jet trigger is an example of a trigger based on the combination of a global variable (E_T-miss) and localized RoIs in the detector. The bulk of the trigger rate will result from fluctuations in the energy measurements of QCD jets, partly as the result of the presence of material in front of the calorimeters and in the regions between the various calorimetry sections. The main instrumental effects arise from the difference in response between the various calorimeter technologies and from the fact that the calorimetry is highly non-compensating.

The contribution of the EF to reducing the LVL1 E_T-miss trigger rate should be important for three main reasons. Firstly, accurate calorimeter calibration and inter-calibration are available [13-1]. Secondly, a separate calibration for cells between clusters can be utilized. Thirdly, the cell E_T cutoff applied to suppress noise can be tuned accurately [13-50]. Initial studies of the E_T-miss + jet trigger rate utilized samples of QCD di-jets (from Pythia) using both fast and full simulation. An excellent agreement was obtained between both simulation methods. Details of the simulation can be found elsewhere [13-51].

The E_T-miss + jet trigger rate calculation has been revisited utilizing the latest ATHENA object oriented reconstruction software framework. In addition, the following improvements were included in the analysis: an updated version of Pythia; an updated model for minimum bias data;

a new algorithm for p_T-miss calculation utilizing the H1 calibration approach as well as tuned cell cuts. Initially the low luminosity scenario has been considered without pile-up and with electronic noise added to the simulation. The preliminary results give a E_T-miss + jet trigger rate of 15/8 Hz for jet-E_T>60/70 GeV and E_T-miss 60/70 GeV. A plot of the integral trigger rate, against the p_T-miss, for both thresholds, is shown in Figure 13-21.

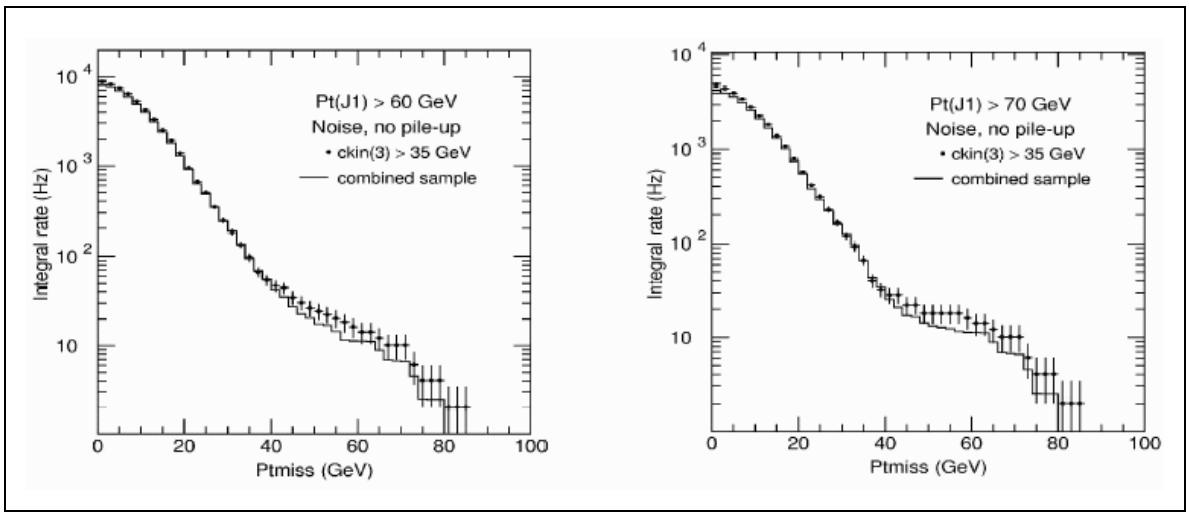

Figure 13-21 The E_T-miss + jet trigger rate, for two P_T (jet) thresholds, as a function of E_T-miss.The plots are for 1×10^{33} cm^{-2} s^{-1} running without pile-up included. Electronic noise (2σ) has been added.

Recently estimates of the tau + E_T-miss HLT rate have been made using the ATHENA reconstruction software. In this case the HLT rate was estimated by assuming that the existing offline tau and E_T-miss code to simulate the performance of the tau + E_T-miss HLT. The effects of electronic noise in the calorimetry were included in the simulation. This preliminary analysis of the low luminosity trigger rates did not include the effects of pileup.The same offline tau identification criteria bulleted in Section 13.4.3.1.3 were used. Also, a cut of ~2σ in the electronic noise was applied to all calorimeter cells immediately after reconstruction and before any further analysis. In this case an improvement was made to the standard offline analysis by making the tau identification criteria dependent on the E_T range [13-52]. In this way the cuts could be optimised for the highest rejection for a given efficiency. It was seen that the jet rejection came mainly from the E_T-miss and after that from the tau identification criteria. This preliminary analysis had only limited statistics and thus only a rough estimate of the trigger rate is possible. However, fixing our cuts to give a E_T tau efficiency of 55% and taking the tau E_T < 100 GeV, the estimated tau + E_T-miss rate is ~5Hz.

13.4.3.3 Triggering on Jets

The purpose of the high level jet trigger is to reduce the rate of events containing jets compared to LVL1 by exploiting the improved E_T measurement. This improvement is achieved by applying a refined energy calibration and jet definition. As jets are the dominant high-p_T process rate reduction cannot be expected by removing fake jet objects. In contrast to e.m. clusters, jets cannot be flagged as background. This would require the reconstruction of exclusive final states which is nor foreseen at LVL2, but would be possible at the Event Filter (EF) level. Jets are searched for in the region of $|\eta| < 3.2$ in the e.m. and hadronic calorimeters. Currently there are three different jet triggers foreseen: j180, j75x3, and j55x4 at low luminosity and j290, j130x3, and j90x4 at high luminosity.

The LVL2 jet finding is guided by the LVL1 regions of interest (RoI). Jets are sought in a region around these RoI's. In the current view the EF performs jet finding across the whole rapidity region with the improved jet calibration and inter-jet calibration allowed at this level. The performance of the high level jet trigger is described in detail elsewhere [13-53]. The jet rates in this report were obtained using fully simulated QCD di-jets.

Table 13-7 The top portion of the table shows LVL1 rates of the different trigger menu items obtained using fully simulated events [13-53]. The middle portion of the table shows the estimated rates using a fast simulation [13-54]. The bottom portion of the table shows the two jet trigger rate for the menu values indicated.

item	1×10^{33} cm^{-2} s^{-1} LVL2+EF (Hz)	Item	1×10^{34} cm^{-2} s^{-1} LVL2+EF (Hz)
j180	253 ± 26	j290	275 ± 87
3j75	286 ± 28	3j130	440 ± 110
4j55	127 ± 15	4j90	165 ± 67
j180	228 ± 8	j290	270 ± 33
3j75	314 ± 10	3j130	270 ± 33
4j55	153 ± 7	4j90	149 ± 25
2j180	101 ± 6	2j290	109 ± 21
2j130	377 ± 11	2j230	358 ± 38

The ATLAS fast simulation program ATLFAST (version 00-02-22) was used to generate 13M di-jet events that were utilized to provide an estimate of the HLT (LVL2+EF) jet trigger rates for one, two, three and four jet final states [13-54]. In this case the jet finding algorithm required a jet initiator E_T of 5 GeV, a jet-cone radius R=0.4 and a jet E_T threshold of 10 GeV. The resulting single and multi-jet rates for low and high luminosity are shown in Figure 13-22. These rates

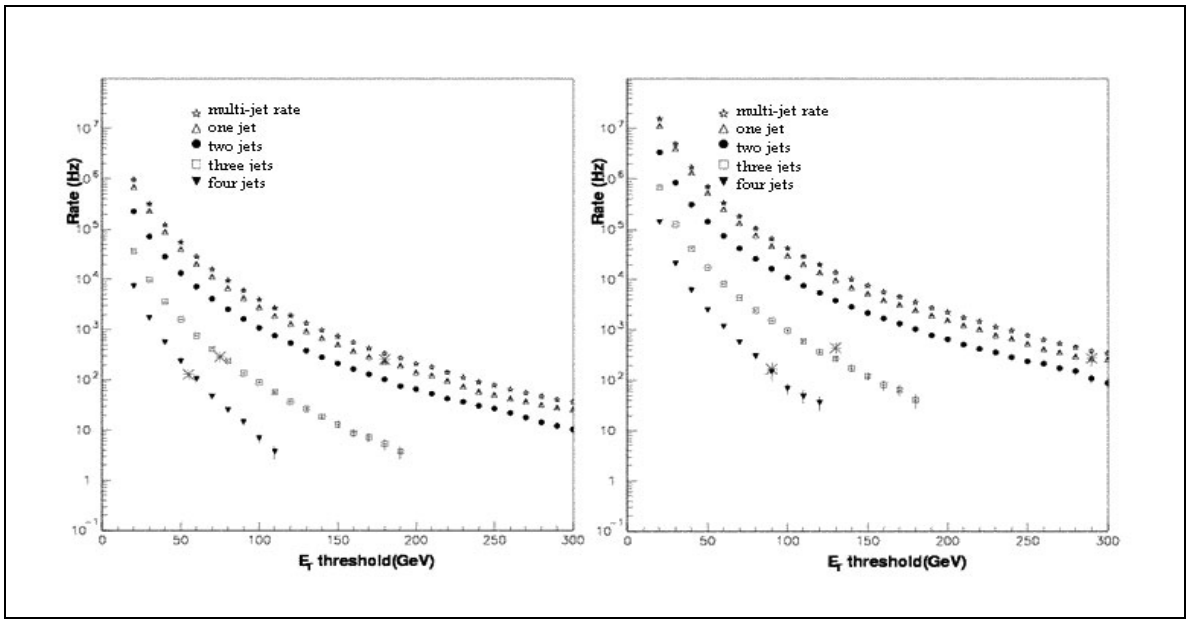

Figure 13-22 The one, two three, four and combined multi-jet rates obtained using ATLFAST. The Large stars show the rates obtained in a previous analysis using full simulation [13-53]. The same scaling factor was applied to both the low and high-luminosity rates obtained from ATLFAST to normalize them to the previous analysis.

were normalized, using a single scaling factor, to the rates predicted, for low and high luminosity running, in a previous analysis based on fully simulated QCD di-jets [13-53]. The one, two, three, and four jet trigger rates for the trigger thresholds given in reference [13-53] and obtained using the ATLFAST simulation, are reported in Table 13-7.

13.4.4 b-tagging

The selection of b-jets at trigger level can improve the flexibility of the HLT scheme and possibly extend its physics performance. In particular, for topologies containing multi b-jets, the ability to separate b-jets from light quark and gluon jets could increase the acceptance for signal events (if the use of lower jet E_T thresholds than those discussed in Section 13.5 is feasible) or reduce the background (and hence the rate) for events containing b-jets that have already been selected by other triggers.

The study presented in this section defines and characterizes, in the low luminosity case, a b-jets selection for the LVL2 based on the information coming from the Inner Detector [13-66]. The extension of the selection to the EF and the definition of an overall strategy for the b-tagging will be adressed in future studies.

13.4.4.1 LVL2 track reconstruction for b-tagging selection

The track reconstruction and the precise determination of the track parameters (in particular the impact parameter (d_0) in the transverse plane) are crucial ingredients of the b-jet trigger.

Several tracking algorithms based on the silicon detectors have been presented in Section 13.3. This study is based on SiTrack [13-65] a LVL2 algorithm that selects and builds track candidates using triplets of space points, because of its good impact parameter resolution.

An early version of SiTrack, PixTrig (based solely on the Pixel Detector), has been used in the Technical Proposal [13-3] to perform a similar study.

The recent change in the Pixel Detector geometry (larger B-layer radius) caused a relevant degradation of the impact parameter resolution for tracks built using only the Pixel Detector. The impact parameter resolution has been recovered and slightly improved at high p_T (Figure 13-23) using the capability of SiTrack to use layer of the Inner Detector: the larger lever arm granted by a SCT layer ensures a better resolution of the track parameters in transverse plane.

13.4.4.2 b-tagging algorithm

The b-tagging selection starts with the track reconstruction performed by SiTrack within the LVL1 jet RoI. For each reconstructed track the significance of the transverse impact parameter $S=d_0/\sigma(d_0)$ is computed; the error on the impact parameter $\sigma(d_0)$ is parametrized, using simulated events, as a function of p_T.

The b-jet estimator is then build using the likelihood-ratio method: for each track (i), the ratio of the probability densities for the track to come from a b-jet or a u-jet is calculated: $f_b(S_i)/f_u(S_i)$; the product W of these ratios over all reconstructed tracks in the jet is computed and the final tagging variable $X=W/(1+W)$ is defined. Jets are tagged as b-jets if X~1 and u-jets if X~0. The selection efficiency of the b-jets and the rejection of light flavour jets can be tuned by cutting on the X variable.

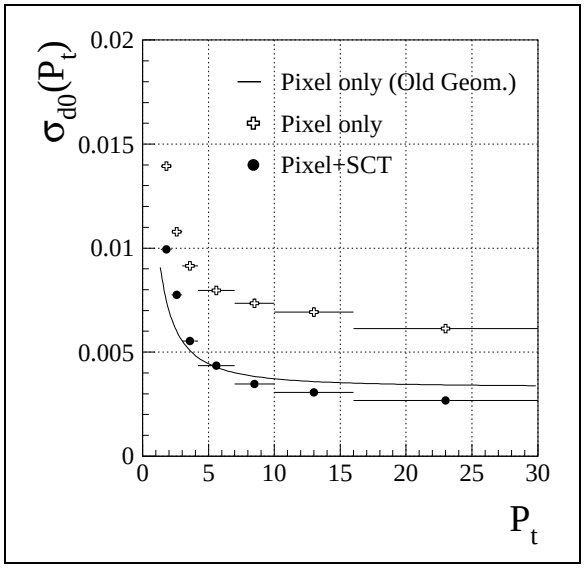

Figure 13-23 Comparison of the d_0 resolution as a function of p_T for different configuration of the LVL2 tracking algorithm SiTrack. The label "Old Geometry" refers to the results obtained in [13-3] with a previous geometry of the Pixel Detector (the so-called Physics TDR layout).

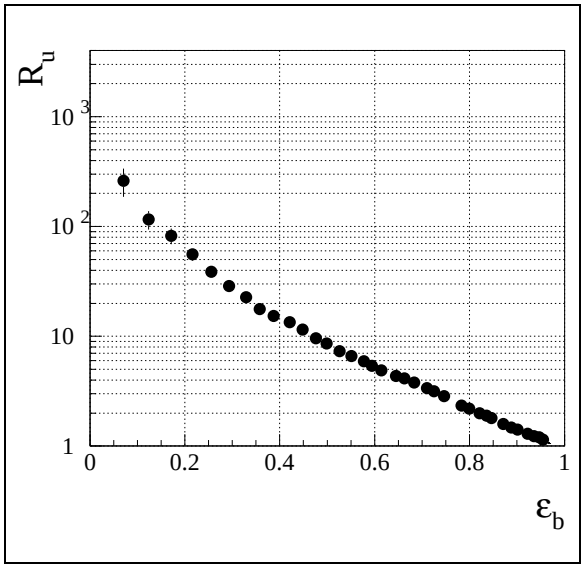

Figure 13-24 u-jet rejection as a function of the b-jet efficiency for jets issued by the decay of Higgs bosons with m_H=120 GeV for Trigger and Offline algorithms. The bias of a Trigger selection on the Offline selection is shown.

The processing time of the b-tagging selection is dominated by the track reconstruction phase and is well compatible with the LVL2 latency (less than 8 ms per jet on 2 GHz CPU on dijet events).

13.4.4.3 Results on single b-jet tagging

The b-tagging algorithm has been characterized on single b-jets coming from H->$b\bar{b}$ decays with m_H=120 GeV produced in association with a W at low luminosity, and corresponding u-jets (taken as representative of the light flavour jets) obtained by artificially replacing the b-jets from the Higgs decay; the LVL1 jet RoI has been simulated selecting a region $\Delta\phi x\Delta\eta$=0.4x0.4 centered around the direction of the quarks issued by the decay of the Higgs boson. The E_T spectrum of these jets covers the range up to E_T=120 GeV providing hence a good benchmark for many physics channels involving Higgs production.

The efficiencies for b-jets (ε_b) and rejection factors (R_u) against u-jets (defined as the inverse of the efficiency for u-jets) are given in Table 13-8. The modest rejections obtained do not spoil the interest of applying a b-tagging selection at LVL2: in multi b-jets events loose selections (necessary to minimize the bias on the selected sample) can still produce significant rejections.

13.4.4.4 Comparison with Offline b-tagging

The performance of the LVL2 trigger algorithm has been compared to that of the Offline algorithm.

The Figure 13-24 demonstrates that the trigger and offline selection are well correlated and that, as long as the LVL2 efficiency is kept above XX%, it is possible to provide subsequent analyses with an unbiased sample in the region ε<XX%.

Table 13-8 Rejection of the LVL2 b-tagging algorithm against u-jets for three different values of the b-jet efficiency: 60%(top), 70 %(middle) and 80 %(bottom) at low luminosity. The results are shown for different intervals of E_T, η of the jet.

	E_T < 40 GeV	40 GeV < E_T < 80 GeV	80 GeV < E_T < 120 GeV
$\|\eta\| < 1.5$	4.0± 0.2	4.6±ð¹±± ± 0.5	5.3 ± 0.6
	2.4± 0.1	3.3 ± 0.3	3.8 ± 0.6
	2.0 ± 0.1	2.3 ± 0.2	2.5 ± 0.3
$\|\eta\| > 1.5$	2.4 ± 0.2	4.5 ± 0.5	2.9 ± 0.6
	1.9 ± 0.1	2.9 ± 0.4	2.2 ± 0.6
	1.4 ± 0.1	2.3 ± 0.3	1.6 ± 0.3

Different combinations of working points of LVL2 trigger selection and offline analysis could be chosen depending on the topology of the events and on the required Offline b-tagging efficiency. Additional flexibility can be added to the b-tagging selection by integrating the LVL2 selection with a selection at EF, where tracking with offline quality will be available.

13.4.5 B-physics

About one collision in every hundred will produce a $b\bar{b}$ quark pair. Therefore, in addition to rejecting non-*bb* events, the B-trigger must have the ability to identify and select those events containing B-decay channels of specific interest. Important areas include CP-violation studies with the channels $B_d \to \pi^+\pi^-$ and $B_d \to J/\psi K_s$ (with both $J/\psi \to e^+e^-$ and $J/\psi \to \mu^+\mu^-$); measurements of B_s oscillations in $B_s \to D_s\pi$ and $B_s \to D_s a_1$ with $D_s \to \phi\pi$; analysis of $B_s \to J/\psi$ and $B \to J/\psi\eta$; rare decays of the type $B_{d,s} \to \mu^+\mu^-(X)$; *b*-production measurements and precision measurements with *B*-hadrons. High statistics are required for these precision measurements. The large number of $b\bar{b}$ pairs produced at the LHC mean that ATLAS is well placed to make a significant contribution in these areas.

Since the Technical Proposal the B-trigger has been re-assessed in the light of a number of developments, including the likelihood of a reduced ID layout at the start of running, an increase in the target start-up luminosity and various trigger deferral scenarios. The aim is to provide the maximum possible coverage of key B-physics channels within the available resources.

It is important to study a range of scenarios since the actual start-up conditions are uncertain, luminosity is expected to vary from fill-to-fill, and there are uncertainties in the physics cross-sections and in the calculation of required resources. A flexible trigger strategy has, therefore, been developed based on a di-muon trigger (with p_T thresholds of about 3-6 GeV) at luminosities of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ or above and introducing other triggers at lower luminosities, either lower luminosity LHC fills or later in the beam coast (over the period of a beam-coast the luminosity is expected to fall by about a factor of two). Two strategies have been investigated for these additional triggers, as follows.

- The baseline strategy is to require at least one LVL1 JET or EM RoI in addition to a single-muon trigger (p_T>8 GeV). At LVL2 and the EF, track reconstruction is performed within RoI using pixel, SCT and, optionally, TRT information. The reconstructed tracks form the basis of selections for e.g. J/ψ(ee), B(ππ) and D_s(φπ). Since track reconstruction is performed inside RoI, the resources required are modest. More details are given in Ref. [13-68].
- If the LVL1 RoI multiplicity proves too high (or the efficiency too low) for the above approach, a fall-back solution is to perform track reconstruction within the full acceptance of the SCT and pixel detectors (full-scan) for events with a single-muon trigger (p_T>8 GeV). In order to minimise execution time, the TRT is not included and so the J/ψ(ee) trigger is not possible. The reconstructed tracks form the basis of selections for e.g. B(ππ) and D_s(φπ). This requires somewhat greater resources than the baseline method, in order to perform the full-scan, but promises better efficiency. This strategy is described in detail in Ref. [13-67].

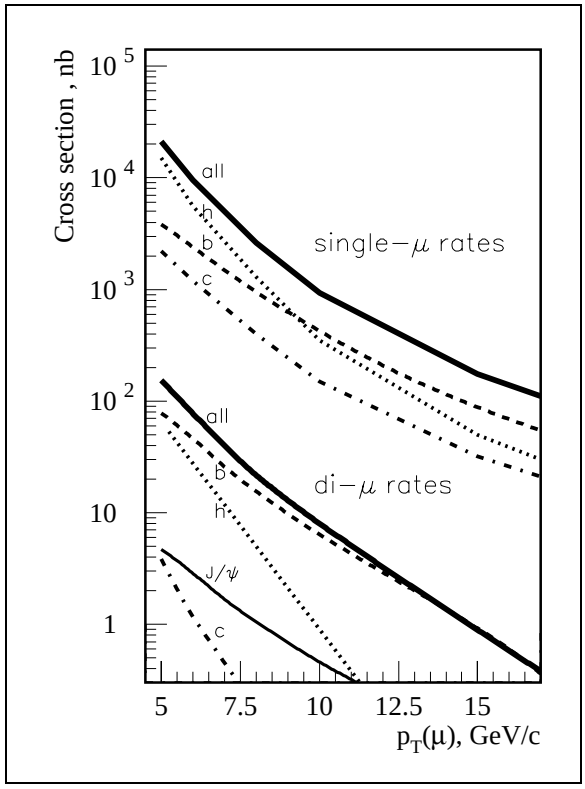

Figure 13-25 Single muon and di-muon cross-sections calculated for a luminosity of 1×10^{33} cm^{-2} s^{-1}. Curves are shown for muons from K and π decays (h), b and c decays and for the sum of these sources (all). The first muon in the event has $|\eta| < 2.4$, the second $|\eta| < 2.5$. For di-muons, the horizontal axis shows the lower p_T. At least one muon must have $p_T > 6$ GeV.

In all cases, at least one level-1 muon trigger is required to initiate the B-trigger. Since the cross-section for inclusive muon production from pion and kaon decays falls more rapidly with p_T than that for prompt muon production from *b*-decays, see Figure 13-25, an appropriate choice of p_T threshold gives a powerful reduction of the trigger rate due to background processes. For example, a threshold of $p_T > 8$ GeV would gives a single-muon trigger rate of 10 kHz at LVL1 for a luminosity of 1×10^{33} cm^{-2} s^{-1}. Most of this rate is due to muons with true p_T below threshold originating from pion and kaon decay, a large proportion of which can be rejected at LVL2 on the basis of more precise track measurements. After the LVL2 selection the trigger rate is about 2 kHz at a luminosity of 1×10^{33} cm^{-2} s^{-1}; about one third of this rate is due to b→μ decays. It is important not to set the muon p_T threshold too high as this would significantly reduce the statistics in the signal channels and render the measurements un-competitive. The rate is further reduced by requiring other triggers in addition to the muon, as described in the following sections.

13.4.5.1 Di-muon triggers

A di-muon trigger provides a very effective selection for several important channels, e.g. B→J/ψ($\mu^+\mu^-$)K_s and B→$\mu^+\mu^-$(X). The LVL1 muon trigger is efficient down to a p_T of about 5 GeV in the barrel region and about 3 GeV in the end-caps. However the actual thresholds used for the di-muon trigger will be determined by rate limitations. For example, a p_T threshold of 6 GeV

would give a di-muon trigger rate of about 600 Hz after LVL1 at a luminosity of 2×10^{33} cm^{-2} s^{-1}. These triggers are mostly due to muons from heavy flavour decays, see Figure 13-25, plus some single muons which are doubly counted due to overlaps in the end-cap trigger chambers. The latter are removed when the muons are subsequently confirmed at LVL2 using information from the muon precision chambers and ID. At the EF, tracks are refit and specific selections made on the basis of mass and decay length cuts. These consist of semi-inclusive selections, for example to select $J/\psi(\mu^+\mu^-)$ decays with a displaced vertex, and in some cases exclusive selections such as for $B_{d,s} \rightarrow \mu^+\mu^-$. Estimated trigger rates are shown in Table 13-9.

13.4.5.2 Hadronic final states

For events with a muon trigger, two strategies have been studied for selecting hadronic final states based on either an ID full-scan or a RoI-based selection. An ID full-scan consists of track-reconstruction within the entire volume of the SCT and Pixel detectors and, optionally, the TRT. The alternative, baseline, strategy uses low E_T LVL1 jet clusters to define RoIs for track reconstruction in the ID. By limiting track reconstruction to the part of the ID lying within the RoI, about 10% on average, there is potential for a significant saving in execution time compared to the full-scan (by up to a factor of ten, depending on the RoI multiplicity per event). Preliminary studies of efficiency and jet-cluster multiplicity have been made using a fast simulation. This uses a rather detailed simulation of the calorimeter which includes a parameterization of longitudinal and transverse shower profiles and takes into account the effects of pulse history, digitization and Bunch Crossing IDentification (BCID). These studies indicate that a threshold of $E_T >$ 5 GeV gives a reasonable mean jet cluster multiplicity of about two RoI per event for events containing a muon with $p_T > 6$ GeV[1], see Figure 13-26.

Table 13-9 Estimated trigger rates for example B-trigger selections at luminosities of 2×10^{33} cm^{-2} s^{-1} and 1×10^{33} cm^{-2} s^{-1}. In these examples the di-muon selection at LVL2 consists only of confirming the muons. The division of selections between LVL2 and the EF remains to be optimised.

	2×10^{33} cm^{-2} s^{-1}		1×10^{33} cm^{-2} s^{-1}	
Trigger	**LVL2**	**EF**	**LVL2**	**EF**
$B_{d,s} \rightarrow \mu^+\mu^-(X)$	200 Hz	small	100 Hz	small
$J/\psi(\mu^+\mu^-)$		10 Hz		5 Hz
$D_s(\phi\pi)$	–	–	60 Hz	9 Hz
$B(\pi\pi)$	–	–	20 Hz	3 Hz
$J/\psi(ee)$	–	–	10 Hz	2 Hz
Total	200 Hz	10 Hz	190 Hz	20 Hz

These studies have been repeated using a full GEANT-based detector simulation followed by a simulation of the LVL1 calorimeter trigger. All sources of electronic noise are simulated, both for the LAr and Tile calorimeters, as well as for the trigger towers. The effect of the BCID system was not simulated, however, and as a result there is a significant contribution to the measured

1. These studies used a muon p_T threshold of 6 GeV. The LVL1 Jet and EM RoI multiplicities can be assumed to be the same for events selected with a 8 GeV muon threshold, within the accuracy of these measurements.

multiplicity from small signals from other bunch crossings. The mean number of Jet RoI found per event is shown in Figure 13-26 as a function of the E_T threshold. The mean multiplicities are significantly higher than obtained with the fast simulation. This is most probably explained by the lack of BCID. Once BCID is included, the measured multiplicities are expected to be much closer to the results of the fast simulation, although this remains to be verified. The full simulation results, therefore, represent an upper limit, with the expected multiplicity lying closer to the fast simulation results. The full simulation shows that a LVL1 jet trigger with an $E_T > 6$ GeV would give an 80% efficiency for finding a RoI containing the signal decay in $B_s \rightarrow D_s\pi^+$ events containing a B_s meson with $p_T > 16$ GeV.

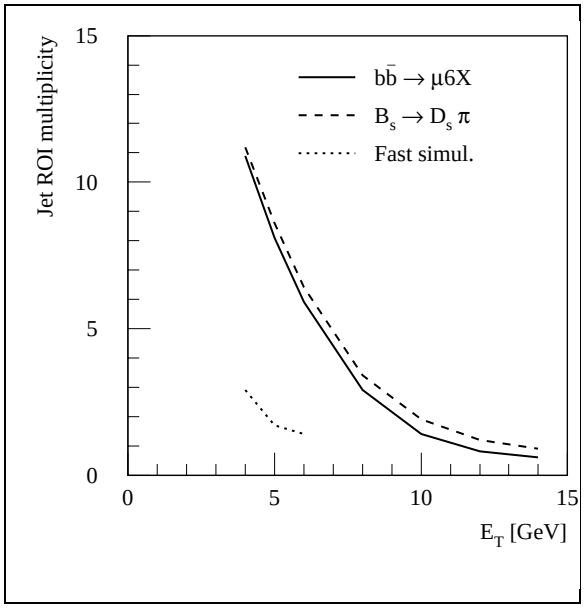

Figure 13-26 The mean number of Jet RoI per event shown as a function of E_T cut. Results are shown for bb→μX events with muon $p_T > 6$ GeV from both the fast simulation (dotted curve) and full simulation (solid line). The mean multiplicity is also shown for $B_s \rightarrow D_s\pi^+$ signal events.

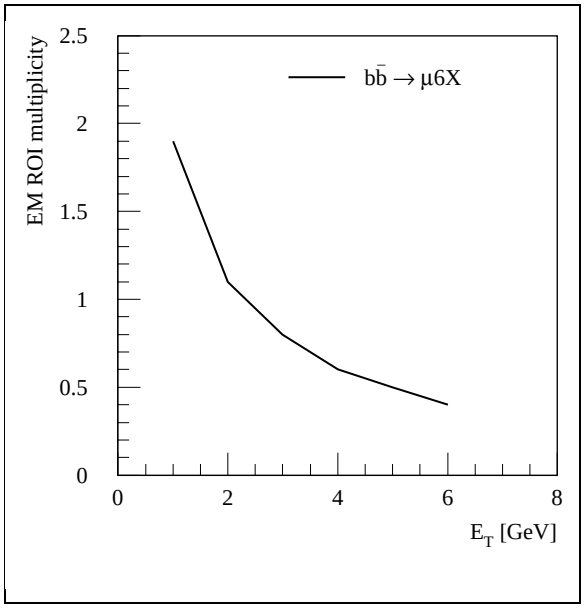

Figure 13-27 The mean number of EM RoI per event shown as a function of E_T cut. Results are shown for the fast simulation for bb→μX events with muon $p_T > 6$ GeV.

Based on the ID tracks reconstructed by one of the above methods (either full-scan or RoI-based) further selections are made for specific channels of interest. These are kept as inclusive as possible at LVL2 with some more exclusive selections at the EF. For example, samples of $B_s \rightarrow D_s\pi$ and $B_s \rightarrow D_s a_1$ events can both be triggered by selecting events containing a $D_s(\phi\pi)$ candidate.

Tracks are refit at the EF inside RoI defined from the results of LVL2. Using LVL2 to guide the EF reconstruction reduces the amount of data to be processed. For example, a region encompassing all LVL2 tracks forming $D_s(\phi\pi)$ or $B(\pi\pi)$ candidates corresponds to about 10% of the ID acceptance, on average. Tighter mass cuts can be applied at the EF than at LVL2 since there is more time available for a more detailed track reconstruction which yields better track parameter resolution. In addition, EF selections may include decay vertex reconstruction, allowing further cuts on vertex-fit quality and decay length.

Studies using a full detector simulation have shown that an efficiency of about 60%[1] can be obtained for $B_s \rightarrow D_s\pi$ signal events where all final state particles have $p_T > 1.5$ GeV. The corresponding trigger rates are shown in Table 13-9; further details are given in Ref. [13-70]. It has

been shown that, with appropriate tuning of selection cuts, there is very little degradation of trigger performance if the pixel layout is changed from three barrel layers to the two layers expected at the start of LHC running. For example, studies based on the IDSCAN algorithm show that reducing the cut on the number of SCT and pixel space-points to a minimum of four for the initial layout (compared to a minimum of five for the full layout) yields the same efficiency as for the full layout with only a 10% increase in trigger rate. Other studies have shown that the trigger is insensitive to the anticipated levels of misalignment in the ID [13-69].

13.4.5.3 Muon-electron final states

A muon-electron trigger is used to select channels such as $B_d \rightarrow J/\psi(e^+e^-)K_s$ events with an opposite side muon tag, or $B_d \rightarrow J/\psi(\mu^+\mu^-)K_s$ with an opposite side electron tag. The LVL1 trigger is used to find low E_T electron/photon clusters which define RoI to be investigated at LVL2. Preliminary studies, using a fast simulation, show that a reasonable compromise between RoI multiplicity and electron efficiency might be obtained with a threshold of E_T>2 GeV. This gives a mean RoI multiplicity of about one for events containing a muon with $p_T > 6$ GeV, see Figure 13-27. These studies give the corresponding efficiency to create a RoI for both the e^+ and e^- from $J/\psi(e^+e^-)$ as about 80% in events where both final state particles have $p_T > 3$ GeV. At LVL2, the electron/photon RoIs are confirmed in the calorimeter, using full-granularity information and including the pre-sampler. Inside the RoI, a search is then made for tracks in the SCT, Pixels and, optionally, the TRT. The RoI about each electron candidate can be quite small, of order $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. This gives a large saving in reconstruction time, compared to a full-scan, but requires a higher p_T threshold than is possible with the ID full-scan. The tracks are refit at the EF, including a vertex fit, and cuts are applied to the decay length and fit quality. The estimated trigger rates are shown in Table 13-9.

13.4.5.4 Resource estimates

In order to estimate the computing resources required for the B-trigger, measurements of algorithm execution time have been combined with estimates of trigger rate at each step of the selection. Various reconstruction algorithms have been timed on several different platforms in order to determine the mean execution time at a given luminosity, the scaling of execution time with the number of hits in an event, and hence the scaling with luminosity. These timing measurements have been combined with the estimates of trigger rates and RoI multiplicity to give an estimate of the resources required for the B-trigger [13-71]. It is estimated that ~15 additional 2 GHz processors are required for the B-trigger, using the baseline RoI guided strategy. This is increased to about 50 for the fall-back full-scan strategy. These estimates are based on execution times measured in a previous trigger performance framework, CTrig. Preliminary measurements in the Athena framework give algorithm execution times (excluding data access overheads) that are a factor 2-3 higher than the CTrig measurements. It is anticipated that with further optimization, algorithm speeds will approach those in the CTrig framework.

In summary, the use of low E_T LVL1 RoI to guide reconstruction at LVL2 promises a full programme of B-physics for very modest resources. However multiplicities and efficiencies need to be verified in studies using a full detector simulation and including a full simulation of the LVL1, with BCID, and higher level triggers.

1. These studies were based on a full-scan of the SCT and pixels at LVL2 with a p_T threshold for reconstructed tracks of 1.5 GeV. The efficiency for a RoI-based trigger has yet to be measured.

13.5 Event rates to off-line

After the submission of the HLT/DAQ Technical Proposal, the baseline scenario for the start-up of the LHC has been changed, now assuming an initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. The information already available from the detailed HLT performance studies, as documented in the HLT/DAQ Technical Proposal [13-3], was used to derive the ATLAS HLT trigger menu for this new scenario, concentrating on the major part, comprised of the mostly inclusive high-p_T object selections. In addition, the impact of constraints coming from limited resource availability was taken into account by restricting the b-hadron physics selection at peak luminosity to di-muon signatures only (for more details see the discussion in the previous section). The resulting trigger menus, both for LVL1 and for the HLT, are presented in Table 13-10 and Table 13-11.

The signatures shown in the trigger menus are mandatory to guarantee an as unbiased as possible and complete coverage of the ATLAS physics programme for observations of new physics and precision measurements. As discussed in Chapter 4, the aim is to be as open as possible to (yet unknown) signatures of new physics, to avoid as much as possible any bias introduced in the trigger selection and to allow for refined selection criteria to be used in the offline analyses. It should be noted that these menus assume stable operation of the LHC machine and the ATLAS detector, and thus do not apply to the initial start-up phase of ATLAS and the LHC, which is addressed in more details in the following section.,

Table 13-10 LVL1 trigger menu with rates for a luminosity of 2×10^{33} cm^{-2} s^{-1}

LVL1 signature	rate (kHz)
EM25I	12.0
2EM15I	4.0
MU20	0.8
2MU6	0.2
J200	0.2
3J90	0.2
4J65	0.2
J60+XE60	0.4
TAU25+XE30	2.0
MU10+EM15I	0.1
others (pre-scaled, exclusive, monitor, calibration)	5.0
Total	~25.0

Table 13-10 lists the major LVL1 selection signatures together with the rates expected for an initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. For most of the signatures, the nominal threshold given corresponds to an efficiency of about 95% for a physics object having a value of the transverse momentum (energy) identical to the threshold value. This is achieved by an appropriate tuning of the selection cuts at LVL1. The selection signatures at LVL1 include single- and di-object isolated electromagnetic calorimeter clusters (which provided candidates for electron and photon object reconstruction at the HLT), single- and di-muon candidates, and the combination of an isolated electromagnetic cluster and a low-p_T muon. In addition, hadronic jets are

selected for various multiplicities (and thresholds). Here it is expected to also include a di-jet trigger (with a threshold lower by several 10 GeV than the one of the inclusive single-jet trigger), which is not shown here. Finally, signatures requiring the presence of missing transverse energy, in association with either a jet or a tau hadron candidate are present. About 20% of the total rate of 25 kHz are allocated for pre-scaled, exclusive, calibration, monitor and other triggers, examples of which are described in Chapter 4**.**

It should be noted that the LVL1 trigger rates are based on the prediction of a physics event Monte Carlo generator and of the detailed simulation of the ATLAS detector response. As such, they do not contain any safety factor against uncertainties on the rates, where possible source are discussed in more detail later. For the design value of the maximum LVL1 accept rate of 100 kHz, this trigger menu thus allows for a safety factor of 4 against possible uncertainties affecting the LVL1 trigger rate.,

Table 13-11 HLT trigger menu with rates for a luminosity of 2×10^{33} cm^{-2} s^{-1}

HLT signature	rate (Hz)
e25i	40
2e15i	<1
γ60i	25
2γ20i	2
μ20i	40
2μ10	10
j400	10
3j165	10
4j110	10
j70+xE70	20
τ35+xE45	5
2μ6 with mass cuts (J/ψ, ψ', B)	10
others (pre-scaled, exclusive, monitor, calibration)	20
Total	~200

Table 13-11 shows the HLT output rate for the above LVL1 trigger menu and does also not include any safety factor against uncertainties on the rates. The rates shown are obtained using the selection algorithm steps at LVL2 and the EF for the various objects, as described in detail in the previous sections of this chapter. The LVL1 electromagnetic selections separate at the HLT into single-/di-electron and into single-/di-photon signatures, which comprise together about 1/3 of the HLT output rate. About 25% of the HLT output rate is used by the single- and di-muon triggers, whereas single- and multi-jet triggers constitute about 15% of the total rate. Selections involving missing transverse energy constitute about 15% of the rate and only 5% of the HLT rate at peak luminosity are allocated to b-hadron physics related triggers (the low-p_T di-muon signature with additional mass cuts to select J/ψ, ψ' and rare B-decays). About 10% of the total rate is allocated for pre-scaled and other triggers.

As discussed in [13-3] and in the previous sections, the ATLAS capabilities and the HLT algorithm performance lead to rather pure samples of inclusive signatures being accepted by the trigger. For the single electron trigger, less than 1/3 of the accepted events contain a fake electron signature and about XX% are due to W→ev decays, where the dominating contribution to the remaining fraction is due to electrons from b/c decays. In case of the di-electron trigger, almost no fake signatures are accepted and the rate is dominated by events containing a Z→ee decay. In case of the single muon trigger, more than XX% of the accepted events contain a muon from a W decay.

SOME MORE NUMBERS/EXAMPLES??

As already mentioned for the LVL1 trigger menu, it is important to note that the rate estimates are the result of simulations, which start with a physics event generator (mostly PYTHIA) and then involve a detailed GEANT based simulation of the ATLAS detector response. They are thus subject to several sources of uncertainties, which are described in the following list:

- **knowledge of cross-sections:** the knowledge of several cross-sections (e.g. the ones for multi-jet production) have sizeable uncertainties (factors of 2 or larger), which could influence the background contributions to the trigger rates,
- **realistic detector behaviour and performance:** the results presented are obtained using the simulated detector behaviour, thus they apply only for stable running conditions with a commissioned and well understood detector,
- **background conditions:** unforeseen changes in beam related background conditions could have an impact on the trigger rates, leading to increased rates (possibly only in certain regions of the detector), and thus the need for a more refined treatment to reject those events efficiently,
- **resources constraints and software performance:** the actual performance of the selection software and the resources available to execute this software might limit the rate rejection capabilities of the HLT (or intermediate stages of it) and thus influence the choice of thresholds for the physics objects used in the on-line selection.

As mentioned already, the above menus are shown for the case of an initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. It is expected that during a machine coast, as the luminosity drops by a factor of 2 or more, or for machine fills which do not reach this peak luminosity value, the selection will be enlarged by assigning a larger fraction of the available bandwidth to pre-scaled triggers (or to include additional exclusive or topological selections, which are not activated at peak luminosity). This would result in the writing of events at a constant rate to mass storage.

13.6 Start-up scenario

The TDAQ system will have to cope, in particular during the start-up phase of the detector operation, with variable data taking environment. It will have to accommodate:

- quickly evolving luminosity of the LHC
- variable beam-related background
- variable electronic noise conditions
- variable configuration of sub-detector electronics, subject to final adjustment iterations

- deferred data handling capacity of the initial TDAQ system
- learning stages of the detector operation

The TDAQ system commissioning has thus to be considered as a self-adjusting process, optimised to use efficiently data available at each stage of the detector operation: at first data collected in sub-detector stand-alone runs, then data collected in the cosmic-muon runs, subsequently data collected in single beam operation of LHC, and eventually beam-beam collision data.

13.6.1 Prerequisites

The commissioning phase of the TDAQ will follow the commissioning phase of the LVL1 trigger and of the sub-detectors. At this stage the following milestones will have to be achieved:

- bunch-crossing synchronization of the RPC, TGC and calorimetric LVL1 triggers at the level of the CTP
- synchronization of the RoI event record at the level of RoIB
- full survey of rates of the LVL1 trigger elements as a function of thresholds, detector region, proton currents, specific luminosity, number of bunches, beam life-time, and background rates (including both the single beam and beam-beam operation phases)
- data produced by sub-detector RODs are debugged
- the background environment of detector operation is well understood from dedicated analysis, at the level of ROD-crate DAQ, of large samples of random events (bunch-bunch interaction timed, pilot bunch timed, and out-of-bunch timed)
- the delay scans for each of the sub-detectors are done using the dedicated, best signal-to-noise, LVL1 triggers and precise timing of detector signals with respect to LVL1 trigger is established
- all methods of preprocessing and formatting data at the ROD level are tested
- the size of background condition and luminosity dependent data volumes are determined for each of the sub-detector RODs

Each sub-detector (sub-detector partition) passing the above ``*ready-for-integration*'' milestones will participate in the subsequent phases of commissioning of the DC, LVL2 and EF systems.

13.6.2 TDAQ commissioning

In the initial phase the rate of the events selected by the LVL1 trigger system will be reduced such that the corresponding data-flow could be comfortably handled by the TDAQ system. The sub-detector ROD data will be transmitted via ROLs to ROBs and made available to the DC system. In this phase the ROL, ROBIn and ROS will be commissioned. In the subsequent phase the ROBIn event-data blocks, bypassing the LVL2 will be transmitted by the EBN to SFIs to be monitored and analyzed.

One of the aims of these runs is to test data traffic and event building capacities of the TDAQ in particular to identify possible sources of traffic jams of the DC system when exposed to increasing LVL1 accept rate and increasing event size. In addition, these runs - taken with dedicated

LVL1 triggers - will be used to collect large samples of ROD-transparent data for fine-tuning of the sub-detector settings and for time-alignment of sub-detector responses.

13.6.2.1 LVL2 trigger and EF commissioning

Once the sub-detector data will be understood, the next commissioning phase will be devoted to debugging and fine-tuning of the LVL2 and EF selection of events. The commissioning of the LVL2 and of the EF will be facilitated by splitting it into several phases. In the first phase the RoIB driven access to ROB data will be tested using single-inclusive LVL1 triggers. In the following one, the LVL2 trigger system will be tested, at first running dummy selection algorithms, then running a simple data validation algorithm for an inclusive trigger and eventually employing the full set of the PESA algorithms. These selection algorithms will be tested in the offline environment using data collected during LVL2/EF transparent running. At first both the LVL2-accepted and LVL2-rejected events will be registered. Once the evaluation of the LVL2 trigger performance will meet the performance targets the LVL2 decisions will become effective.

Following the above phase, the next one will be devoted to testing of the DFMs and SFI performance and subsequently of the performance of the EF selection of events. At this stage, combined studies of the performance of the data selection and of the data collection capacity will be made. They will include studying the saturation effects with increasing LVL1 accept rate, and with increasing event sizes. These studies will allow to assess the capacity of the actual (deferred) TDAQ system including detailed evaluation of the data collection, data preparation, event selection and event building performance of TDAQ for various LVL1-types of events.The results of these studies, together with the results of the survey of the LVL1 trigger rate, will allow to determine the luminosity-phase dependent running strategy of the ATLAS detector. This strategy will be based on the optimal use of the resources available at the start-up of the detector operation and on the best use of luminosity available for physics runs.

13.7 References

13-1 ATLAS detector and physics performance technical design report, CERN–LHCC/99–14/15 (1999)

13-2 ATLAS Collaboration, First-Level Trigger Technical Design Report, CERN/LHCC/98-14.

13-3 ATLAS HLT, DAQ and DCS Technical Proposal, CERN-LHCC/2000-17

13-4 T. Schörner-Sadenius and T. Wengler, Formats and Interfaces in the Simulation of the ATLAS First Level Trigger, ATL-DA-ES-0029.

13-5 M. Abolins et al., Specification of the LVL1 / LVL2 trigger interface, ATL-DA-ES-0003.

13-6 T. Schoerner-Sadenius, The ATLAS Level-1 Trigger offline Simulation: Overview, Core PArts, and Use, ATL-COM-DAQ-2003-008.

13-7 E. Moyse and A. Watson, Performance and Validation of TrigT1Calo, the offline Level-1 Calorimeter Trigger Simulation, ATL-COM-DAQ-2003-010

13-8 http://xml.apache.org/xerces-c

13-9 E. Eisenhandler, Level-1 Calorimeter Trigger URD, ATL-DA-ES-0001.

13-10 A. Watson, Updates to the Level-1 e/gamma and tau/hadron Algorithms, ATL-DAQ-2000-046.

13-11 ATLAS Collaboration, Muon Spectrometer TDR, CERN/LHCC/97-22.

13-12 A. Di Mattia and L. Luminari, Performance of the Level-1 Trigger System in the ATLAS Muon Spectrometer Barrel, ATL-DAQ-2002-008
http://atlas.web.cern.ch/Atlas/GROUPS/MUON/layout/muon_layout_P.html

13-13 A. Di Mattia et al., A muon trigger algorithm for Level-2 feature extraction, ATL-DAQ-2000-036;

13-14 A. Di Mattia, RPC Trigger Robustness: Status Report, ATL-DAQ-2002-015.

13-15 S. Veneziano, Preliminary Design Report of the MUCTPI Electronics, ATC-RD-ER-0001.

13-16 N. Ellis, Central Trigger Processor URD, ATL-DA-ES-0003; P. Farthouat, CTP Technical Specification, ATL-DA-ES-0006.

13-17 G. Schuler et al., Central Trigger Processor Demonstrator (CTPD), ATL-DA-ER-0005; I. Brawn et al., A Demonstrator for the ATLAS Level-1 Central Trigger Processor, ATL-DA-ER-0008.

13-18 R. Blair et al., ATLAS Second Level Trigger Prototype RoI Builder, ATL-D-ES-0011.

13-19 S. Armstrong et al., “Requirements for an Inner Detector Event Data Model”, ATLAS-TDAQ-2002-011.

13-20 A.G. Mello, S. Armstrong, and S. Brandt, “Region-of-Interest Selection for ATLAS High Level Trigger and offline Software Environments”, ATLAS-COM-TDAQ-2003-005, ATLAS-COM-SOFT-2003-002.

13-21 PESA Core Algorithms Group, S. Armstrong editor, “Algorithms for the ATLAS High Level Trigger”, ATLAS-COM-TDAQ-2003-00X.

13-22 K. Assamagan et al., “A Hierarchical Software Identifier Scheme,” ATLAS-COM-MUON-2002-019.

13-23 C. Leggett and A. Schaffer, presentations at the ATLAS EDM-DD Workshop, 27 January 2003.

13-24 For more information on SCTKalman see I. Gaines, S. Gonzalez and S. Qian, in Proceedings of CHEP2000 (Padova); D. Candlin, R. Candlin and S. Qian, in Proceedings of CHEP01 (Beijing); J. Baines, et al. ATL-DAQ-2000-031.

13-25 J. Baines et al., “Global Pattern Recognition in the TRT for B-Physics in the ATLAS Trigger”, ATLAS-TDAQ-99-012; M. Sessler and M. Smizanska, “Global Pattern Recognition in the TRT for the ATLAS LVL2 Trigger”, ATLAS-TDAQ-98-120.

13-26 S. Sivoklokov, presentations made in PESA Core Algorithms Group meetings, December 2002 and January 2003. See also S. Sivoklokov, “High pT Level-2 Trigger Algortihm for the TRT Detector in ATRIG”, ATLAS-TDAQ-2000-043.

13-27 M.P. Casado, S. González, and T. Shears, TrigT2Calo package, http://atlas-sw.cern.ch/cgi-bin/cvsweb.cgi/offline/Trigger/TrigAlgorithms/TrigT2Calo/

13-28 S. González, T. Hansl-Kozanecka, and M. Wielers, “Selection of high-pT electromagnetic clusters by the level-2 trigger of ATLAS,” ATLAS-TDAQ-2000-002.

13-29 S. González, B. González Pineiro, and T. Shears, “First implementation of calorimeter FEX algorithms in the LVL2 reference software,” ATLAS-TDAQ-2000-020.

13-30 S. González and T. Shears “Further studies and optimization of the level-2 trigger electron/photon FEX algorithm,” ATLAS-TDAQ-2000-042.

13-31 ATLAS first level trigger technical design report, CERN-LHCC/98-14 (1998).

13-32 H. Drevermann and N. Konstantinidis, “Determination of the z position of primary interactions in ATLAS,” ATLAS-TDAQ-2002-014.

13-33 H. Drevermann and N. Konstantinidis, “Hit FIltering and Grouping for Track Reconstruction,” ATLAS-TDAQ-2003-XXX (Document in Preparation).

13-34 A. Di Mattia et al. (Rome Level-2 Muon Trigger Group), “A Muon Trigger Algorithm for Level-2 Feature Extraction,” ATLAS-DAQ-2000-036.

13-35 I. Gavrilenko, “Description of Global Pattern Recognition Program (xKalman)”, ATLAS-INDET-97-165; also see: http://maupiti.lbl.gov/atlas/xkal/xkalmanpp/index.en.html.

13-36 R. Clifft and A. Poppleton, IPATREC: Inner Detector Pattern Recognition and Track Fitting, ATLAS-SOFT-94-009.

13-37 K. Assamagan et al., "Definition of Raw Data Objects for the MDT Chambers of the Muon Spectrometer", ATL-COM-MUON-2003-020

13-38 T.A.M. Wijnen, "The MROD Data Format and the towering partioning of the MDT chambers", ATL-COM-MUON-2002-011.

13-39 P. Bagnaia, A. Di Mattia, S. Falciano, T.A.M. Wijnen, "Online versus offline identifiers for the MDT chambers of the Muon Spectrometer", ATL-COM-MUON-2003-017.

13-40 K. Assamagan et al., "Raw Data Object Definition for the RPC chambers of the ATLAS Muon Spectrometer", ATL-COM-MUON-2003-19

13-41 G. Aielli et al.," Data Format of the RPC Detector of the Muon System", ATLAS-COM-MUON-2003-018.

13-42 J. Shank et al., "Track Reconstruction in the ATLAS Muon Spectrometer with MOORE", ATL-COM-MUON-2003-012, ATL-COM-SOFT-2003-008.

13-43 J.Shank et al.,``Moore as Event Filter in the ATLAS High Level Trigger'', ATL-COM-MUON-2003-013, ATL-COM-SOFT-2003-009,ATL-COM-DAQ-2003-012.

13-44 B. Caron, J. L. Pinfold, R. Soluk, “A Study of the ATLAS Tau/hadron Trigger”, ATL-COM-DAQ-2003-030, June 2003

13-45 ATLAS Trigger Performance Status Report, CERN/LHCC 98-15, 25th August 1998

13-46 A. T. Watson, “Updates to the Level-1 e/gamma & tau/hadron Algorithms,” ATL-DAQ-2000-046.

13-47 B. Caron, J. L. Pinfold and R. Soluk, “The ATLAS Tau/hadron Trigger at LVLE and LVL2”, ATL-DAQ-2003-XXX

13-48 B. G. Piniero, “Tau Identification in the Second Level Trigger”, ATL-DAQ-1998-127

13-49 D. Cavalli and S. Resconi, “Combined Analysis of $A/H \rightarrow \tau\tau$ Events from Direct and Associated bbA Production,” ATL-PHYS-2000-005.
D. Cavalli, L. Cozzi, L. Perini, S. Resconi, “Search for $A/H \rightarrow \tau\tau$ Decays”, ATL-PHYS-94-051
D. Cavalli and S. Resconi, “Tau Jet Separation in the ATLAS Detector”, ATLAS PHYS-N0-118

13-50 D. Cavalli, “Missing Transverse Momentum Reconstruction in ATLAS”, ATL-PHYS-1996-080

13-51 G.Polesello, D. Cavalli, B. Caron, R. A. Davis, “Event Filter Rates for the E_T-miss + jet Trigger at Low Luminosity”, ATL-DAQ-2000-016, March 2000.

13-52 M. Heldmann, D. Cavalli, S. Resconi, “Status Report on the τ-Identification Performance”, ATLAS Physics Workshop, Athens, June 2003.

13-53 M. Wielers, “Performance Studies of Jets in the High Level Trigger”, ATL-DAQ-2000-015, May 25th 2000.

13-54 J. de Jong and J. l. Pinfold, “Estimating the ATLAS Multi-jet Trigger Rate Using ATLFAST”, ATL-COM-DAQ-2003-029.

13-55 Performance studies for electron and photon selection at the event filter, ATLAS internal note ATL-DAQ-2000-007 (2000)

13-56 First study of the LVL2-EF boundary in the high-p_T e/gamma high-level-trigger, ATLAS internal note, ATL-DAQ-2000-045 (2000)

13-57 Electron trigger performance studies for the event filter, ATLAS internal note, ATL-DAQ-2001-004 (2001)

13-58 J. Baines et al., “Performance Studies of the High Level Electron Trigger”, ATL-COM-DAQ-2003-0020 (2003)

13-59 T. Sjostrand, Comput. Phys. commun. 82 (1994) (1994) 74; T. Sjostrand et al., Comput. Phys. Commun. 135 (2001) 238

13-60 E. Richter-Was, “Samples from the DC1 production: dataset 2000, dataset 2001, dataset 2003; few items from physics content evaluation“, ATL-COM-PHYS-2003-026, (2003)

13-61 E. J. W. Moyse, A. Watson, “Performance and Validation of TrigT1Calo, the offline Level-1 Calorimeter Trigger Simulation”, ATL-COM-DAQ-2003-010 (2003)

13-62 ATLAS Data Challenge 1, see http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/DC/DC1/DC1-Report-V1.0-211002.pdf (will become ATL-SOFT note)

13-63 Di Mattia A., Falciano S., Nisati A., "The implementation of the muFast algorithm in the new PESA framework", ATL-COM-DAQ-2003-024.

13-64 A. Virchaux et al., "MUONBOX: a full 3D Tracking Programme for Muon Reconstruction in the ATLAS Spectrometer", ATL-MUON-1997-198.

13-65 M.Cervetto et al., SiTrack: a LVL2 track reconstruction algorithm based on Silicon detectors, ATLAS-COM-DAQ-2003-35

13-66 M. Cervetto et al., b-tagging Event Selection for the ATLAS High Level Trigger: an update. ATLAS-COM-DAQ-2003-XX

13-67 J.Baines et. al., B-Physics Event Selection for the ATLAS High Level Trigger, ATLAS Note ATL-DAQ-2000-031 (2000).

13-68 A.Watson, V.Ghete, A study of the use of Low E_T calorimeter RoI in the ATLAS B-Trigger, to be written.

13-69 *Effects of Inner Detector Misalignment and Inefficiency on the ATLAS B-physics Trigger* by: J.Baines, B. Epp, S.George, V.M.Ghete, G.Hollyman, D.Hutchcroft, A.Nairz, S.Sivoklokov. ATL-DAQ-2001-006

13-70 *Event Filter Rate for the Ds Trigger* B. Epp, V.M.Ghete, A.Nairz. ATL-DAQ-2001-003

13-71 J.Baines et. al. Resource Estimates for the ATLAS B-physics Trigger ATLAS-COM-DAQ-2002-001

13-72 N.Konstantinidis and H.Dreverman, Fast tracking in hadron collider experiments, in Proceedings of the 7th International Workshop on Advanced Computing and Analysis Techniques in Physics Research, Amer. Inst. Phys. Conference Proceedings, Vol. 583, 2001

13-73 J. Baines et al., Pattern Recognition in the TRT for the ATLAS B-Physics Trigger, ATLAS Note ATL-DAQ-99-007 (1999).

13-74 P.Billoir and S.Qian, Simultaneous Pattern Recognition and Track Fitting by the Kalman Filtering Method, Nucl. Instr. and Meth. A225 (1990) 219.